Notes on building a transistorised, but otherwise 'classical', Termenvox (“Theramin”)

Jamie Davies, **MM0JMI**

jamie.davies@ed.ac.uk

This diagram is to show general layout: the circuits are shown at higher resolution later in the text.

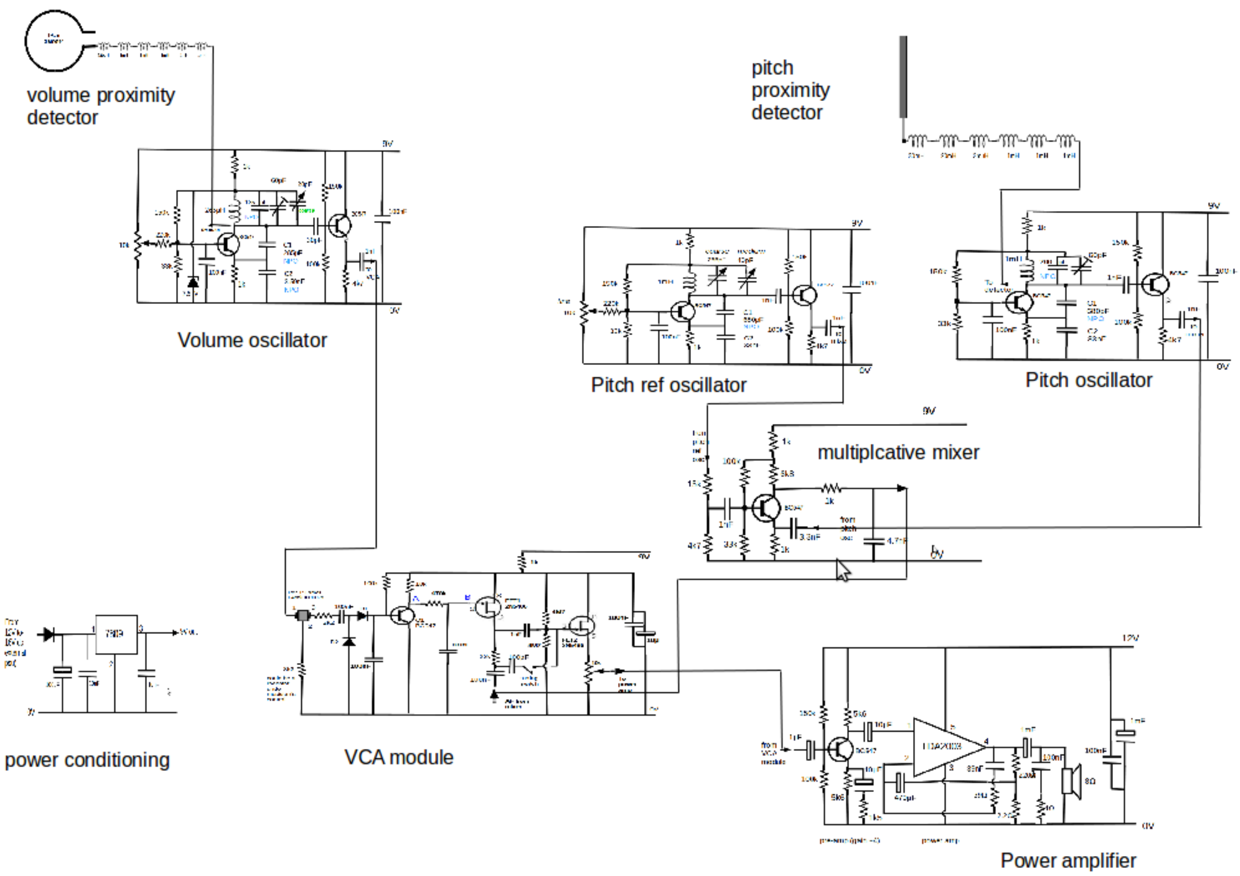

Contents:

1. Introduction to this document

This document describes the construction of a an electronic instrument based closely on the Etherphone/ Termenvox/ Theramin of the 1920s and 1930s. The way the instrument works, and the way in which it is played, are intended to be faithful to the original but the original thermatron* circuits have been replaced by transistor versions, allowing the instrument to be smaller, more portable and safer.

The design presented here is by no means all my own. I began by building an instrument exactly as described by Barry Holloway on his excellent page www.strangeapparatus.com. In my hands, at least, this design resulted in a device that worked well enough to demonstrate how a Termenvox functions. It fell short, however, of what I consider to be the requirements of a real musical instrument to be played by a real musician (especially a musician who has little patience for 'character' in machines and just wants things to work). This comment is in no way meant to be a criticism of Barry – he may well be a highly skilled builder who managed to achieve more stabilty and linearity from his design than I could with my more modest abilities. To get the behaviour I wanted, especially stability in low pitches, large dynamic range in volume, and reliability across a range of temperatures and humidities, I had to make a number of changes. These resulted in a machine that is only very slightly more complex in terms of component count, but that is more suited to musical performance in venues varying from hot cabaret clubs to very cold cathedrals.

In these notes, I have tried to explain how everything works, and especially to explain the reasoning behind specific choices of design. This helps readers debug their circuits, and perhaps to make different decisions. In a number of places in the text, I have mentioned broad RF design principles that would need much more space to explain than is reasonable to take up in this document: to help readers who want to learn more, and who do not want to take my word for something, I have included references to the kinds of books common on the shelves of electronics enthusiasts and libraries. These are cited by an abbreviated version of the title, and the references are listed in full at the very end of the document.

The text assumes that a reader is reasonably familiar with building electronics, but not necessarily with analogue RF circuits. Readers who are familiar with standard RF buidling blocks may therefore find that some sections are over-explained: just skip those sections.

Disclaimer:

The instrument was built by, and this document has been written by an amateur enthusiast. No warranty is implied. Also – to state the obvious – your safety that that of others around you, as you work with tools and electricity, is your own responsibilty, not the author's.

---

* ie UK 'valve' / US 'tube': the international term 'thermatron' has been introduced by Grayson Evans TA2ZGE in his excellent book on contemporaty uses for these devices, *Hollow State Design*.

2. The history of the original Etherphone/ Termenvox/ Theraminvox/ Theramin

The Termenvox was probably the world's first electronic musical instrument. It was invented by the Russian engineer Ле́в Серге́евич Терме́н ('Lev Sergeyevich Termen') in 1920. Терме́н did not set out to create a musical instrument: his invention illustrates the role of serendipity in technology.

Терме́н was employed to measure the dielectric constants of gases in the Soviet Union's brand new Physical Technical Institute. To make these measurements, Терме́н, already experienced in radio work, developed a device of great sensitivity. He built two high-frequency oscillators whose tuning network included an air-gap capacitor. One of these oscillators was in a standard gas (eg dry air) and the other in the test gas. If the dielectric constant of the test gas differed from that of the standard gas, the capacitance would change, and the test oscillator's frequency would differ by a small amount from that of its standard twin. The signals from two oscillators were mixed together in a way that created the sum of the two original frequencies and one at the difference of them. The difference frequency appeared in the audible range (100-20,000Hz) and could be heard through headphones or a loudspeaker, and could be used to estimate the change in oscillator frequency.

The story goes that, when moving his hand near the oscillators, Терме́н heard the difference frequency change. He then made a version of the circuit that had nothing to do with test gases, but which had half an air gap capacitor (ie just one plate) attached to an auxilliary resonant circuit in the test oscillator and was sensitive to extra capacitance from a nearby hand, which acted as the other plate. He also added a volume control, initially in the form of a foot pedal. He demonstrated the instrument to his colleagues in October 1920, playing tunes he had learned for his cello. Over the next month, he added an extra oscillator and voltage-sensitive amplifier, so that the volume could be controlled by the proximity of his other hand to a second half-capacitor. This meant that the instrument could be played with no contact at all. In November 1920, he gave his first concert. At around the same time (sources differ on the exact order of events), Терме́н invented a related device, used to detect the proximity of a person for the purposes of acting as a burglar alarm.

Терме́н called his instrument the 'Etherphone'. Within the still-young Soviet Union, where the invention was welcomed as a symbol of technical originality, it became known as the Терменвокс ('Termenvox'). In 1925, Терме́н travelled to Germany to licence his patents to Golberg & Sons: they called the instrument the 'Theraminvox'. In 1927, Терме́н travelled to the USA, where he went by the name Léon Theremin. In that country, his instrument was simply known by his name – the Theramin. While in America, Терме́н/Theramin gained considerable publicity by playing his instrument with the New York Philharmonic, and licenced his design to RCA. He continued design work on other instruments and worked closely with Clara Rockmore, another Russian emigrée who has come to be regarded as the finest ever player of the RCA Theramin and a special version built just for her (see YouTube). Терме́н left the USA for the Soviet Union in 1938 and, after a period in the Gulag after Stalinist purges, and another being involved with electronic aids for espionage, he returned to electronic instruments in the 1960s. He died in 1993.

The instrument, with its unworldly sound, has appeared in the indicental music of many films including *The Day the Earth Stood Still*, *Spellbound* and *The Lost Weekend*. It did <u>not</u> appear, though, in the Beachboys' *Good Vibrations*, which many people cite as the most famous example of its use: the instrument used there achievend a similar sound but, unlike the Termenvox, needed physical contact with the player.

3. Do you really want to build a Termenvox?

Electronic challenges: I do not want to put anyone off enjoying what I have found to be an interesting, though often maddening, project but it is only fair to point out that building a Termenvox is a lot less trivial than its simple-looking circuits might imply. The fundamental operating principles, especially the proximity detectors, are quite subtle and require patience to get right. In a recent conversation on a web forum[*], two skilled builders of electronic instruments, at least one of whom is world-famous in the field, made the following remarks:

Writer 1: *If the Theremin didn't already exist I wouldn't believe that it could exist, that sensing tiny changes in the tiny capacitance of the hand a significant distance away from a small antenna could be stable enough and noise free enough to play serious music. Which I suppose is what keeps me fascinated with it. That, and it's a wonderful playground on which to hone almost everything one knows about electronics.*

Writer 2: *I have never undertaken anything which has educated me more - Back in the 80's I took on a design to control the deposition of a one molecule layer of phospholipids onto a substrate - it had to maintain precisely one layer of lipids on the surface of a trough while the substrate was dipped... I never imagined that I would find something as "simple" as a theremin could pose any challenge after having achieved that! - But that job took less than a year - and here I am after having spent at least 5x as long on theremin R+D. But I have learned more about electronics in this time than over the entire 30+ years preceding it.*

People like this do not remain fascinated for years with tasks that are easy: be prepared to be patient and to learn, even if you are copying the exact design presented here (because so much about the tuning of the proximity detectors depends on exact details of the wiring, the box etc).

Getting this instrument working will be much easier with good test gear. A high-impedance voltmeter is a must (digital is good). I found a frequency counter to be really helpful, and an oscilloscope to be useful, although Терме́н must have managed without them. I also found an audio-frequency signal generator valuable in testing the audio amplifiers, though any audio source (eg via a jack plug from an MP3 player) would probably do. A grid dip meter that works in LF and MF ranges could save you a lot of time in setting up the proximity detectors (mine works only on HF/VHF, so I did not take this short-cut. On reflection, it would have been quicker to build one than to manage without).

Musical challenges: When you have built a working instrument, using it to produce real music (rather than sci-fi special effects) is not at all easy: for all its apparent simplicity, the Termenvox is a very demanding instrument. It requires the player to have excellent relative pitch. Players of bowed instruments (violin, cello etc) will have the required pitch discrimination, whereas players of keyed instruments may not. The musician also needs an ability to hold his or her hands still, or to make small defined movements, and also to hold the head and body still, which is something very few musicians do. As Clara Rockmore observed, this is also the only instrument for which the player must even play the rests (with a normal instrument, a rest means you do nothing: here, it requires a positive movement of the left hand into the volume loop to silence the tone). Playing from memory or by improvisation is almost essential, unless you create some system for making

* http://www.thereminworld.com/Forums/T/28810/antannae-non-linear-relationship?Page=1: read 4 March 2015.

music visible without anything moving near the antennae to turn the pages. Finally, opportunities for playing a Termenvox with a conventional orchestra are very few.

Learning curve challenges: just to state the obvious, buidling a working Termenvox will not mean that you will be playing like Clara Rockmore in a week. Or a year. Getting the best sounds out of a Termenvox requires effort by both the player and the builder: the danger point (especially if the builder and player are two different people) will be when something will not sound 'right' and player and builder blame one another. It is perfectly true that a professional violinist will sound better on a Stradivarius than on a primary school music class violin. It is equally true that a Stradivarius gains its magic only when in the hands of a very skilled player – a beginner could make even this sound awful. If it is really not clear whether the problem is in the instrument or player, try to get the player to a known-good instrument or get the instrument to an expert player. Whatever happens, try not to let something that is supposed to be fun become a source of argument and frustration.

Cost: The circuits described here use mostly standard components of the kind that most electronics enthusiasts will have in bulk. It is very difficult to work out component costs for one project when economical sources of resistors, for example, sell them for a fraction of a penny each, in rolls of a hundred.

Of the components that are *not* just pennies each, the approximate costs are:

- The 19 inch rack case (about £30). There are cheaper, plastic, alternatives.
- Mint tins (about £1.20 each, when bought full of mints)
- Inductors (about £2 each: you will probably need to buy a much larger selection than you end up using, or buy cores and be prepared to mess about taking windings on and off)
- Potentiometers (about £1.50 each)
- Regulator and audio amplifer (about £1 each)
- Ceramic resonator (about £3).
- Polyvaricons (about £3 each).
- Circuit boards (about £1 each when cut down from larger ones).
- Power supply – if you do not have a suitable surplus 'wall wart', you can make a power supply for about £15.

Provided you already have tools and test equipment, the build is therefore inexpensive. If you were starting from nothing, and have to buy all tools and test gear as well as the components of the project, then we would be talking many hundreds of pounds. But you would only be in this position if you did not already make electronic devices. If you are new to electronics, put the Termenvox idea on hold and start by building simpler things to gain experience and confidence, or find a really good, patient Elmer* to help you.

* Elmer: radio ham slang for a gnarled and ancient enthusiast who has been playing with electronics long enough to have made all the mistakes and learned from them, and is willing to pass on Tribal Knowledge – the stuff that never appears in textbooks – to newcomers. Elmers can be recognized easily at ham radio club meetings – they are the ones with the wild-staring eyes and solder burns in their beards (female versions do not always have the beard). The mention of radio ham clubs is appropriate here – the difficult parts of the Termenvox are the parts that operate at <u>radio</u> frequencies, and the Tribal Knowledge for radio frequency work will not be known by people who only deal with audio systems or with digital devices. Indeed, they may do exactly the wrong thing, as they are used to suppressing radio-frequency oscillations rather than encouraging them.

4. The build itself.

4.1 A high-level description of how the whole instrument works.

The cicuitry of the Termenvox described here can be divided, functionally, into nine separate modules (ten if one counts the external DC power supply, but this can be taken as read for now). These modules are: the internal power conditioning unit, the pitch proximity detector, the pitch variable oscillator, the pitch reference oscillator, the multiplicative mixer, the volume proximity detector, the volume variable oscillator, the frequency-controlled amplifier (itself consisting of 3 submodules), and the audio power amplifier. They are arranged as in figure 1:

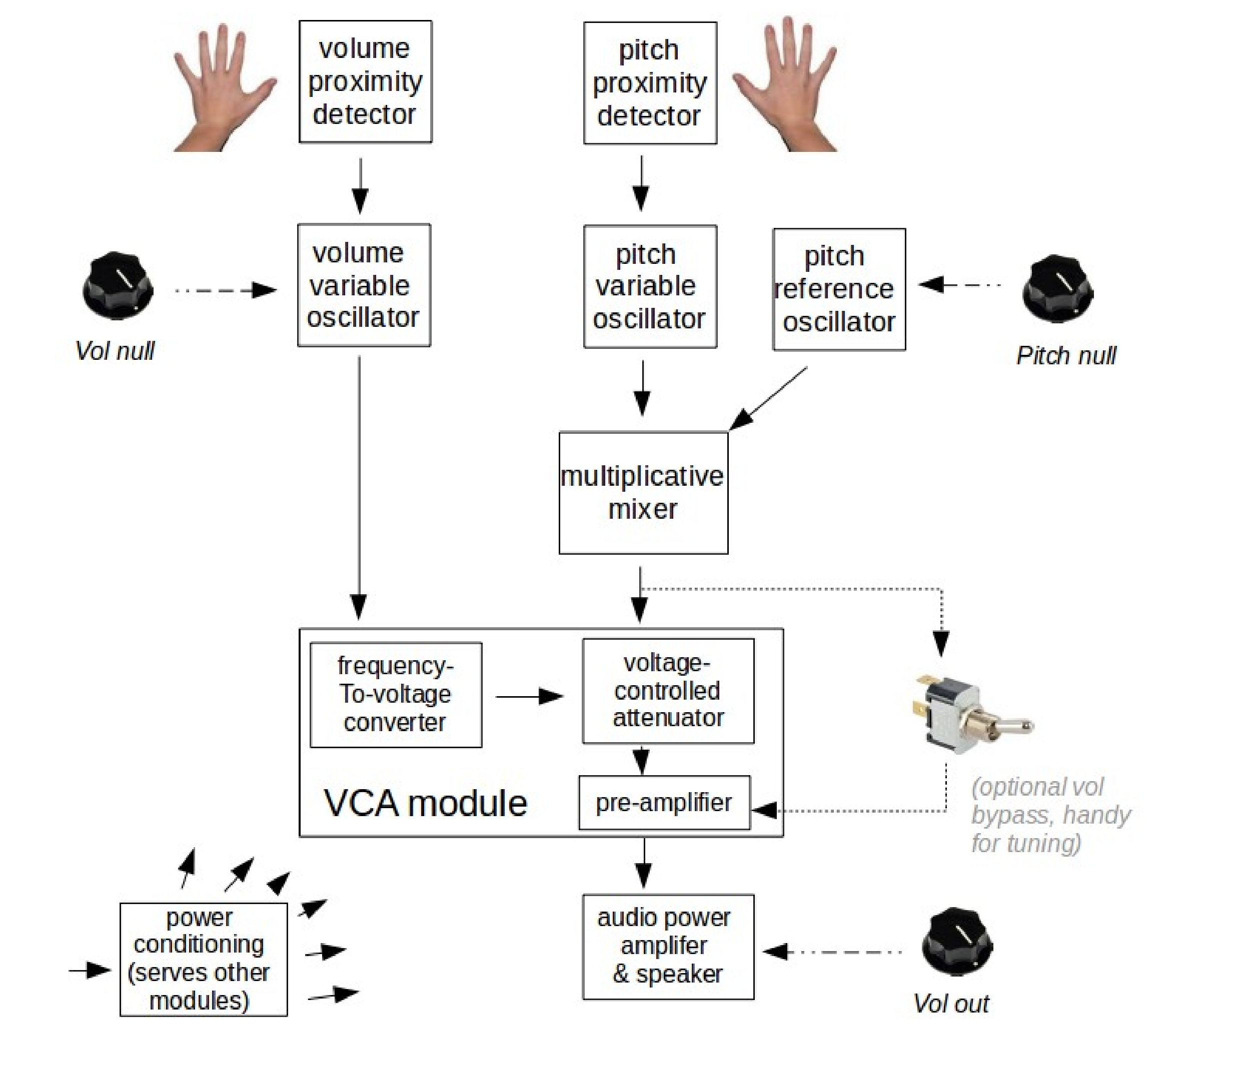

Figure 4.1.1: an overview of the Termen-vox. The 'detectors' at the top are arguably part of the oscillator modules below them: they are being drawn separately here because many designs on the web ignore some critical points about their function, resulting in instruments that require the hands to operate over such a short and such a non-linear range that the device is useless for music. Drawing them as separate functional units helps to focus attention on their design.

The archictecture of the Termen-vox arises from its being played without any direct physical contact between the player and the device. Pitch and volume are controlled by the proximity of one of the players' hands to two half-capacitors (often misleadingly called 'antennae'). The hand acts as the missing half of the capacitor, the resulting capacitance being highest when the hand is closest. The difference between a hand being present and not present is, however, tiny: of the order of one picofarad (1pF, ie 10^{-12}F). Audio-frequency tuned circuits typically involve capacitances of the order of tens of microfarads (10μF, ie 10^{-5}F). Adding or subtracting 1pF to 10μF would make a difference of just 0.000001%, so directly tuning an audio-frequency oscillator over the musical frequency range of 100Hz – 3000Hz by using hand capacitance is clearly out of the question. Radio frequency tuned circuits operating at frequencies of a few hundred kHz typically involve capacitances of the order of magnitude of the 100pF: adding up to 1pF would therefore make a difference of around 1% to the capacitance. Making a 1% difference directly to the capacitance of the main tuning circuit of a radio-frequency (RF) oscillator operating at a few hundred kHz would still have too little effect on frequnecy (which, in most circuits, follows the square root of capacitance). Терме́н realized that it is possible to add a second tuned network to an oscillator and, as long as the main tuned network will sustain oscillation by the usual mechanisms of positive feedback, there is no need to design any feedback into the second one. This gives the designer an unusual amount of freedom in planning the second network, which is called the 'proximity detector' in this document. He can, for example, achieve resonance with high inductance and very low capacitance. Adding a hand capacitance to this causes a great relative change to the total capacitance, and pulls the resonant frequency of this second network strongly downwards; strongly enough that the frequency of the oscillator to which it is connected can drop by several percent. If the no-hands frequency of the oscillator is 200kHz, a change of 0-4kHz (2%) is easily enough to accommodate the required range of audio outputs of 100Hz-3kHz.

The pitch proximity detector, and the pitch variable oscillator with which it interacts, are set to run at the same frequency when there is no hand near the pitch half-capacitor. Most of the setting is done at the end of the construction phase, but final setting of the null difference is done by the musician using the 'pitch null' knobs on the front panel, immediately prior to performance. This pair of knobs (one coarse, one fine) is provided because different environments can cause small variations in the no-hands frequency of the pitch variable oscillator. The RF signals from the two oscillators are combined in the multiplicative mixing module. When a multiplicative mixer is fed with signals at two different frequencies, it creates additional signals at the sum and difference of the original frequencies. The sum signal will be medium-frequency RF, and is discarded. The difference signal (the heterodyne) will vary between zero Hz, when the two oscillators are running at the same frequnecy, to about 3000 Hz when the frequency of the pitch variable oscillator is reduced by the presence of a hand near the pitch half-capacitor. This is the audio-frequency (AF) signal that is amplified and fed to the speaker.

If there were no requirement for volume control, this AF signal could be passed directly to a pre-amplifer, a power-amplifier and the speaker (some 'toy' devices advertised on the web as 'pitch only theramins' work exactly this way). For full Termenvox there is a requirement for volume to be controlled by the closeness of a hand to the volume proximity detector. The volume proximity detector and volume variable oscillator work in exactly the same way as their pitch equivalents, except that the oscillator operates at a higher no-hands frequency of 454kHz, which drops when a hand is brought near. This frequency is again finely adjustable by the musician using the volume null knobs, again to compensate for the environment. In the design described here (which is not the same as Терме́н's) , the signal from the volume variable oscillator is brought to a tuned filter, which is designed to allow signals very near 455kHz to pass through it almost undiminished, but to reduce signals more and more the further their frequencies are from 455kHz. The strength of the signal emerging from the filter is therefore high under no-hands conditions, but much lower when a hand is brought near enough to the volume proximity detector to cause the frequencty of the volume variable oscillator to fall. The emerging signal is turned into smooth dc, and this dc is used to regulate a voltage-controlled attenuator that is situated between the multiplicative mixer and the pre-amplifier. When the voltage is high (no hand present), the attenuator is set to minimum attenuation and most of the AF signal from the mixer is allowed to proceed to the preamplifer. As the control voltage reduces, because a hand comes near the volume proximity detector, attenuation of the AF signal increases, allowing less to proceed to the preamplifer. The volume from the final speaker is therefore less, down to inaudible when a hand is very near the detector.

The final module is a completely standard off-the-shelf audio amplifer: it is equipped with a normal volume knob that sets the maximum loudness of the instrument just as a volume knob may control the loudness of a radio or TV. Not shown in the diagram is an output jack that can allow the musician to play through an external amplifier for performance purposes.

4.2 Design challenges for an instrument-quality Termenvox.

As there is a big difference between a plastic toyshop trumpet and a real musical instrument, so there are differences between 'toy' termenvox devices intended to demonstrate general principles, and perhaps make special effects noises, and ones that can be used as mucical instruments.

The specific challenges in making an instrument not a toy are:

- Building proximity detectors that respond when the hand is a reasonable distance away (1ft/ 30cm) and, in the case of the pitch one, give a response that a musican perceives as being reasonably linear with respect to final musical pitch, over a few relevant octaves.
- Building the pitch control system so that the oscillators can go down to the lowest pitches of the musical range without suddenly 'locking' together into silence.
- Building the volume control system to be sensitive enough to give at least a 20dB (ie 100x power level) difference between 'mute' and 'loud'. I estimated the size of this range by measuring the dynamic range on Clara Rockmore's recording of The Swan.
- Building the volume control system to be fast enough for staccato playing.
- Making the musician's manual controls sensitive enough that they allow a null to be reached, but not so sensitive enough that they are hard to set in the right place.
- Building everything with enough stability that it will tolerate changes of temperature and some degree of physical stress, for example during transport.

4.3 An overview of the construction methods used in this project.

The construction methods used here are very common in amateur radio. Completely different methods could be used but, since some aspects of physical layout are important to function, it is probably a good idea for me to explain the features that were the result of deliberate decisions rather than just a random aspect of how things turned out.

Overall box size: this is a musical instrument played with both hands, which need to be held away from the body. The half-capacitors that project from the main box therefore need to be around shoulder-width apart, so the most natural size for the main box of the project is about shoulder-width. The obvious choice is a 19inch rack box, commonly used for scientific instruments and high-end hi-fi units. These are metal (strong, and an effective RF shield) and, while nobody would call them pretty, they look reasonably professional.

Fig 4.3.1: the finished instrumnent, built into a 19” rack cabinet. It is placed on an organ stool for photography, but would be played in open space.

Space: given the size of the box, there is no point in making the circuits physically small. Small circuits mean having components close together, increasing the risk of unwanted capacitive or inductive coupling. Also, making measurements from small circuits can be frustrating, with a risk of probe tips causing damaging short-circuits. I therefore used the whole space available.

External power supply: I used an external 240Vac → 12Vdc 'wall wart' adaptor rather than building the power supply into the main box for three reasons. The first reason was weight: for safety reasons, given the exposed half-capactiors, you will want to use a transformer-type supply rather than a switched mode unit, and even small (1A) transformers are quite heavy objects. This Termenvox was designed to mount on a camera-type tripod at its middle, so that the forward leg of the tripod is between the feet of the player rather than in the way of either foot. Having a heavy lump inside the box, probably near its edge, would tend to make the tripod want to 'sag' one way. The second reason was a combination of weight and aesthetics: mains ac power cables are thicker, less flexible and more obvious than more lightly insulated 12V supply lines, so having a 12v line running between the Termenvox and the floor is less ugly than having a mains cable. Constructors who wish to hide the supply lines in an elegant stand (for example, connecting the main box to the stand via a large and strong plug) will have an even stronger reason for choosing to use a 12V supply. The third reason was safety during construction: it is much nicer, when making adjustments to a live project, not to have to keep worrying about the risk of touching pieces of metal at 240V ac. I should stress that this particular 'wall wart', which was quite old, was linear (transformer, rectifier, capacitor and 7812 regulator), so safe and free from switching noise. I would advise avoiding switched-mode power supplies, which are likely to introduce all kinds of electronic noise into the system and which have modes of failure that argue against their use for any project that includes touchable metal circuit components.

Modular construction: when building electronic projects, I think in functional modules (like the ones in Fig 4.1.1) and tend to build in this way too. This has the advantages that individual boards are simple and do only a limited number of things, that testing is easy, module-by-module, and that if one wants to make major changes to one module there is a much reduced risk of damaging others in the process. Construction is also more rewarding, as one can build one module at a time, test it, and move on in the certainty that progress has been made. This means that I used seven individual circuit boards, plus small tag-strips for the proximity detector modules. The one disadvantage of modular construction is some inefficiency in component use: the independence of each module reduces opportunities for one component performing multiple functions. It is often said that an engineer is someone who can do for one dollar what any fool can do for two. Insistence on modularity makes my designs veer more towards the 'any fool' component count.

Mint tins: The Termenvox has three RF oscillators that need to be prevented from interacting with one another except at the two places at which they are meant to be brought together (mixer and voltage-controlled amplifier). In particular, the two pitch oscillators must be prevented from interacting because, if they are allowed to do so, there is a very high risk that they will 'grab' each other as their frequencies converge, and lock on to precisely the same frequency. If this happens, low AF pitches will be impossible to maintain. One of the best ways of isolating RF modules is to surround them with metal, and 'standard operating procedure' for this within the amateur radio community is to place each module in a small metal tin. In the UK, the favourite is the *Marks and Spencer Curiously Strong Mints* tin. Putting all modules in tins also has the advantage of protecting the ones on which you are not working, under closed lids, from dropped tools, blobs of solder, flying ends of cut wire etc.

Fig 4.3.2: the use of mint tins to contain the seven main modules of the Termenvox. These are all glued to the inside of what will be the top of the insrument. This photograph was taken before the final changes were made to the design, and in particular the shaft of the coarse pitch control is not present.

Plain wire interconnections (not shielded coax): various designs on the web recommend the use of shielded co-axial cable to connect different parts of (modules of, in my terms) the circuit. There are three things I dislike about this idea. The first and most important is that shielded coaxial cables have transmission line impedances of between 50 and about 600 ohms (a few very expensive ones go a little higher). Every time a signal meets a sudden transition of transmission line impedance, it undergoes reflection and having incoming and reflected signals bouncing about near non-linear components such as transistors is a recipe for the generation of unwanted harmonics. Plain wires tend to have much higher impedences. The second reason is that the exposed braid of the coax needs to be soldered very carefully and completely to avoid the risk of little flecks of metal braid detaching and coming to rest where they are not wanted: these fragments ('devil's whiskers') seem to be peculiarly attracted to places where they can make an almost invisible but highly destructive short circuit. The third reason is that shielded cables for RF use are stiff and get in the way. With modules shielded from interconnecting cables by the mint tins, there seems to be no point in using shielded cable, so I did not. I did, though, try to route different signal cables away from each other as much as possible. In practice, there has been no problem of signal interaction through signal cable cross-talk.

Twisted pair power lines: this is pretty standard practice but worth a mention anyway – twist the positive and negative power supply lines together. If they do pick up stray RF, the two lines will then pick it up equally so that there is no RF voltage between them, and they do not therefore superimpose an RF signal on their dc power. This is one defence against inter-module interaction via power lines. The other is the use of decoupling capacitors across the power lines on each circuit board, which is a standard thing to do.

Long-plastic-shaft potentiometers: these are a good way ot preventing hand capacitance on the tuning knob from getting anywhere near the oscillators, and allow the modules to be mounted well back from the front panel (the instrument is meant to be sensitive to hands near the half-capacitors – sensitivity to hands anywhere else will be a nuisance to the musician).

Manhattan/ ugly-style circuit boards: This aspect of construction will be completely familiar to readers from the amateur radio community but may be quite alien to everyone else unless they started playing with electronics in valve (tube) days. Outside radio, most electronics amateurs seem to like to emulate as far as possible the electronics industry's use of printed circuit boards, in which components are located on one side of a board and interconnections are made by copper tracks on the other, small holes allowing the component leads through. For mass production by unskilled workers or robots, printed circuit boards make a great deal of sense: they can be produced in high volume, drilled in jigs, soldered in one operation in a solder bath, and there is no expectation that they will be altered. For one-off projects in amateur workshops, they make no sense at all: designing and etching a circuit board takes a great deal of time, drilling all of the holes is tedious, soldering has to be done connection-by-connection and when the components are upside-down and falling back out of the board. When the printed circuit is made, the front of the board does not reveal the connections behind it, making measurement and debugging unnecessarily difficult, the board has to be mounted using 'stand-offs', and it has to be unmounted again for any work to be done. Worst of all, changing components is difficult and making large-scale alterations to the design is almost impossible unless one resorts to the 'ugly' methods being recommended here from the beginning.

The popular alternative to a custom-made PCB. 'Veroboard' (holes drilled at 0.1inch spacing, and parallel copper tracks connecting them behind the board) saves work of etching and drilling, but it requires a lot of track breaking for removing unwanted connections, it suffers from horrible capacitive coupling between parallel tracks, and has every other vice of printed circuit boards.

A completely different technique is to build everything on the *front* of an unetched copper-clad board[*]. Ground connections are made directly to the board and oconnections either to component leads ('ugly' construction: QRPB p116-117) or to small pads of PCB material superglued to the top of the board ('Manhattan' construction: QRPB p121-123). These pads can be bought in bulk: I use the ones from *qrpme.com*. Pads can also be home-made with a hole punch spare PCB material. In general, I start with Manhattan construction but some changes are made the 'ugly' way: particularly if I need to add a small capacitor in parallel with an existing large one, I will typically just solder its leads directly to the leads of the existing one. The advantages of Manhattan/ ugly construction are that it is quick (no etching, drilling or cutting), everything, components and connections, is on the same side of the board (no stand-offs, no hidden connections – you can just glue the board into the tin); there is a massive and continous ground-plane, great for suppressing RF interference; connections are easy to trace and can have probes clipped to them, and adding, removing and replacing components is straightforward. Even major changes can be made by gluing on a few more pads. The only disadvantage is that the circuits are aesthetically messy, but who cares? The electrons don't and, once the lid of the mint tin is closed, nobody has to look at what is inside.

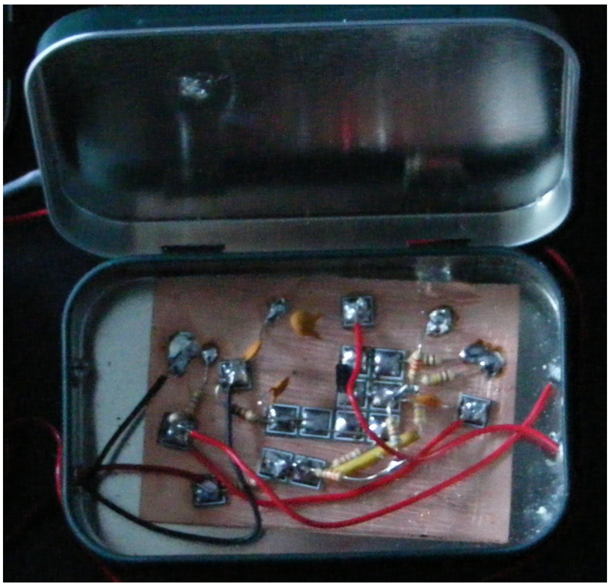

Fig 4.3.3a: Manhattan-style construction of one of the simpler modules (the mixer). The small tinned squares in rows are the Manhatan pads (the square block look is the origin of the name).

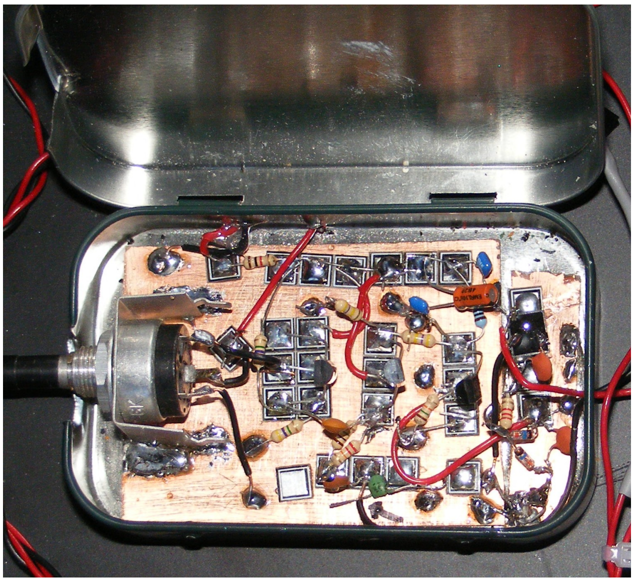

Fig 4.3.3b:The most complex of the modules (the voltage-controlled amplifier).

---

* see EMRFD chapter 1 for more discussion – this reference is written out in full at the end of this document.

One final thing – whichever method of construction is used, the copper of the board will need to be cleaned before use. Fine sandpaper us much better than wire wool, because sandpaper does not shed small pieces of conductive material that will stick to sleeves and tools to be dropped somewhere important and inaccessible.

Order of construction/ testing: the most logical order will depend on the test gear available. Having a working amplifier is useful for testing earlier stages, so if a signal generator is available, I suggest starting with the power conditioning unit, speaker and power amplifer, and working 'backwards' towards the proximity detectors.

4.4 The box, panel, holes etc.

Whatever box you use, drill all of the holes you will need and use a vacuum cleaner to remove every trace of swarf (metal shavings and cuttings) *before* you start adding any electronics. Vacuum the cuffs of your sleeves, too.

Front panel: The Termenvox described here has six manual controls in the front panel. Three of these, the tuning switch and the fine pitch controls, are not strictly necessary although I find the switch really useful. You may also want to consider building in a test tone generator (see section 7): this will need yet another switch. I used Letraset white letters to label the controls. If, as I did, you use the top plate of the box as the chassis and use solder-down type potentiometers (see fig 4.3.3b),, the shafts of the potentiometers will be 'upside-down' so that the screw of screw-type knobs will be visible from the top. Aesthetic musican-types may object to this ugliness: you can file a flat on the other side of the potentiometer shaft to allow the knob to be fitted screw-down.

Rear panel: On the back panel, you will need a socket for dc power (of a design that does not allow a connection to be made wih the wrong polarity), and probably an output socket to feed an external amplifier for performance situations. Make this panel neat, because it will be the audience-side in a performance situation.

Top and sides: *Important* – what follows is correct for a right-handed musician. If the person for whom you are building this is left-handed, reverse the layout.

On what will be the left-hand side of the instrument, from a player's point of view, you will need a socket or sockets for the volume loop half-capacitor. I used two banana plug sockets, centres 1.5 inches apart, one red and one black. One is for an electrical connection; the other is present for purely mechanical reasons. The colours match those on banana plugs on the half-capacator 'loop', and are to ensure that the loop is plugged in the same way each time to avoid one possible source of tuning variability.

The pitch half-capacitor will need to come out of the top of the instrument near its right side. Again, I used a black banana socket – this half-capacitor is a vertical rod, so only one socket is needed. I chose to make the top plate of my box the 'chassis' on to which all of the mint tins are attached (meaning they are all 'upside-down' in use). I decided to do this so that the connection between the volume half-capacitor would be fixed and rigid, as any flapping flexible connection (needed if the bottom plate were the chassis and the top were the 'lid') would risk problems in tuning and stability. But... this decision made final adjustment of the pitch section difficult as the instument had to be propped up in an odd position so that both half-capacitors were in free space but I could still reach the pitch reference module. What might have helped, had I thought of it, would have been to take the connection to the pitch half-capacitor out of the right hand *side* of the box by another pair of bananna plugs (one electrical, the other mechanical). This way, the top of the box would have had nothing coming through it and could have been a simple lid. Doh!

Bottom panel: no cables need to come through this (unless you are planning a clever arrangement in which a heavy-duty socket connects your instrument to an elegant stand and also brings in power and takes out audio to a performance amplifier). The bottom panel probably does need rubber feet top be bolted or glued to its corners (this is important: the corners of 19 inch rack panels are vicious to the surfaces of tables and forgetting rubber feet may get you into serious trouble with the 'domestic authorities'). If (like me) you want to the unit to be mountable on a heavy-duty camera tripod, you will need to arrange an appropriate hole and bolt in the bottom (I MIG-welded a bolt on the inside, in a position well away from the modules. If you do not have a MIG welder, an auto-repair shop will probably be willing to do this for you in a matter of seconds for a contribution to their tea-and-biscuits fund,. But be nice and clean off the paint around the area first. Also, take along a bolt that fits the nut, as it is much easier to weld a nut to a plate when a bolt is already holding it in place. Do not, by the way, try to stick-weld to a plate as thin as that of a 19-inch rack box: you will blow right through the metal.

Edges and corners: unless you are the musician who will be playing this, it is probable that the musician is not used to 'engineering' type things and will not expect to find sharp corners on musical instruments. It is worth going round the box with a file and making sure that everything is smooth enough for hands not to be cut, then blowing some matt/satin black spray paint over the whole thing, which you will probably have to do after drilling anyway. I used high-temperature stove paint just because it was to hand.

4.5 The modules in detail.

4.5.1: The power-conditioning module.

This Termenvox is intended to be run from an external power supply, giving smoothed dc of 12-13.8V at up to 1 Amp. This external supply should be of the old-fashioned transformer type, not a switched-mode unit. There are two reasons for this. The first is that switched-mode power supplies, especially cheap ones, can generate horrendous RF interference that could interact with signals in the Termenvox to create unexpected frequencies. The second is that switched-mode power supplies can suffer modes of failure that connect mains voltages to what should be low-voltage dc: rare as such failures are, they would be dangerous in the context of an instrument that has metal projections that the player can touch. Old-fashioned transformers provide proper isolation. By the way, if you are going to use a surplus power supply ftom a no-longer-wanted item of consumer electronics, as well as opening it up to make sure it really is of the transformer type, do check the output voltage before using it. I have two that both say they give “12V dc 1 Amp”, and both actually give 19V dc at all current drains from 0 to 1A. I added a 7812 regulator (on a heatsink) to one of the 19V PSUs mentioned above to make a 12V one, which was less expensive than buying a new transformer would have been. The unit can also run from a 12V battery – although a lot has been written about the need to ensure mains earthing of a Termenvox to complete the loop between human hand and instrument ground, in practice the unavoidable capactive coupling between the box, stand, floor etc are easily good enough and there is really no need to fuss about additional earths.

The power amplifer can run directly from 12-18V. All other units are run from a regulated 9V supply, in the power-conditioning module. This regulated supply is utterly standard (Fig 4.5.1.1) and uses a 7809 regulator chip. The power drawn off by the modules run from the 9V supply is very modest indeed (about 60mA) so there is no need for a heatsink. There are no special requirements for construction, apart form the obvious one of allowing plenty of space for connecting five positive supply lines to the output side. If Manhattan construction is used, the negative lines can be soldered anywhere on the board. The only test to be done is verifying the 9V.

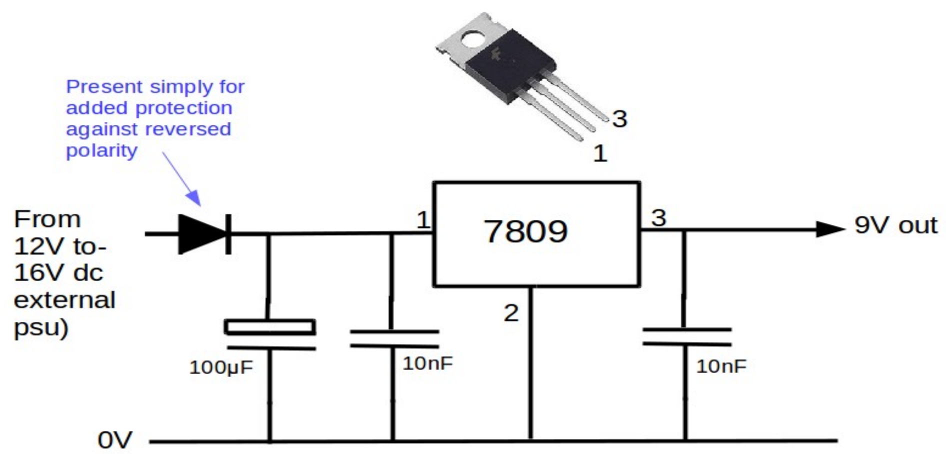

Fig 4.5.1.1: the power conditioning module. The 100µF electrolytic should be rated at 35V or above (electrolytics fail a lot – it is sensible to overspecify). This assumes the external PSU already has a large smoothing capacitor in it. It must, for the sake of the audio amplifier.

4.5.2: The audio amplfier.

The idea of using a TDA2003 amplifer-in-a-chip comes from that of Barry Holloway's article. There were some variations between Barry's circuit diagram and that given by the manufacturers in their data sheet, particularly in relation to the polarity of some electrolytic capacitors (C1, C2). I followed the circuit of the manufacturers' data sheet. I also added a preamplier, to produce the circuit in Fig 4.5.2.1. The gain of the preamplifier will be about 4.5. The estimation of this gain figure can be justified by noting that a transistor will always keep its emitter voltage about 0.6V below its base voltage. If the base voltage swings by δV, the emitter voltage must also change by δV, so by Ohm's law, the emitter current must change by $\delta I = \delta V/R_E$, where R_E is the resistance on the emitter line. Most of this current must flow through the collector, so the collector line experiences the same δI. Again by Ohm's law, the swing in its voltage $\delta V_c = -\delta I\ R_c$, where R_c is the resistance on the collector line. Therefore, substituting for δI, $\delta V_c = \delta V.R_c/R_E$. Gain is output signal/ input signal, = $\delta V_c / \delta V = R_c/R_E$. In this case, the resistance in the collector line is 5600Ω and the resistance to *signals* in the emitter line is about 1200Ω, (because the capacitor lets signals go via the 1500Ω path was well as the 5600Ω dc path; 1500Ω in parallel with 5600Ω gives an effective resistance of 1183Ω, which we may as well round up to 1200Ω). The final gain figure is therefore 5600/1200 ≈ 4.5.

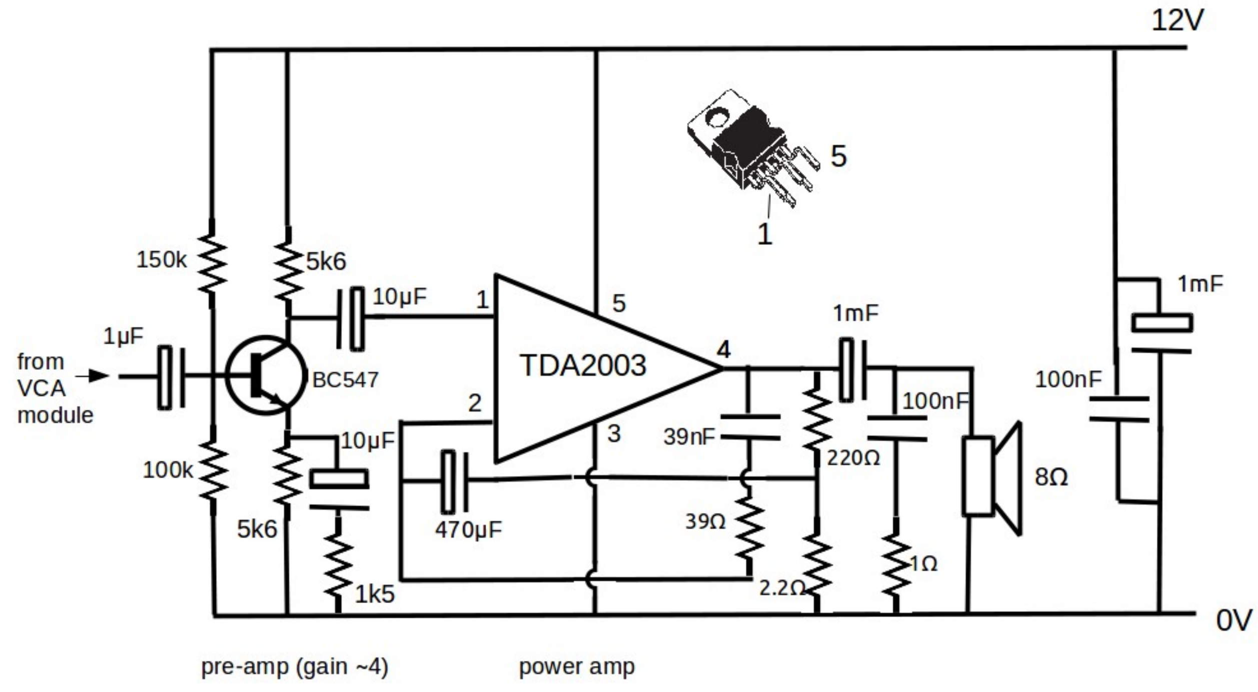

Fig 4.5.2.1: The audio power amplifer. This design appears on the original TDA2003 data sheet from ST microelectronics and is reproduced under the terms of their fair use licence – the whole datasheet is available from www.datasheetcatalog.com

I used a 8Ω loudspeaker, obtained via ebay (nothing special). The metal tab of the amplifier chip is electrically isolated, so I used the copper board itself as a heatsink. In this application, it will be running at low power, and is anyway equipped with intermal thermal protection.

The first test is to power on the amplifer and check that there is no unwanted signal. The most likely one to be met is motorboating ('bub, bub, bub, bub, bub...') somewhere around 1-4 Hz. In general, motorboating has two casuses: lower than expected supply voltage, or poor decoupling. When I built this unit, it started motorboating. Adding a 1mF (ie 1000μF) electrolytic capacitor across the supply lines as physically close to the TDA2003 chip as possible cured the problem completely (I kept the existing 100μF capacitor there too, mainly because I could not be bothered to remove it). I suspect that leading the speaker return directly back to the supply and not to the amplifer board may have been another way to cure the problem.

The next test is to connect an AF signal generator (running at 10-50mV) to the input and check that the amplifer amplifies and that the speaker works.

4.5.3 The frequency-controlled amplifer module.

This module consists of three-submodules, all in the same mint tin. Having them in the same tin makes sense as external connections between them would be at risk of collecting unwanted signals.

4.5.3.1: The preamplifier sub-module.

The point of the pre-amplifier is to present a high impedance to the incoming audio signal (ie to be controlled by the signal's voltage without drawing much current from it) and to amplify it and pass it on at a lower impedance as required by the power amplifier. Having a high input impedance at the preamplifier is useful because the presence of the amplifier has a negligible effect on the performance of the attenuator upstream: apart from anything else, this makes calculations about the attenuator easier to make.

There are 3 basic configurations of bipolar ('normal') transitor amplifiers; common emitter, common collector ('emitter follower') and common base. Of these, the common collector configuration has by far the highest input impedance and by far the lowest output impedance (see *DTFCM* p66): it is effectively a current amplifier that does not amplify the voltage much, and this is just what we need. Field effect transistors (FETs) offer very high input impedances anyway, so using an FET in a common drain configuation (the FET equivalent of the common collector configuration) is ideal for our purposes. The 10k potentiometer between the source terminal of FET2 and ground provides an overall manual volume control for the speaker volume (to avoid annoying the neighbours etc). The 4M7 resistors provide bias to half the effective supply voltage (this is the normal way to bias an amplifier as it accommodates the maximum possible range of input swing without the clipping that would occur if the voltage on the Gate approached that on the other terminals). The 1nF capacitor upstream of the gate isolates the dc bias of the Gate from dc voltages the other side of the capacitor.

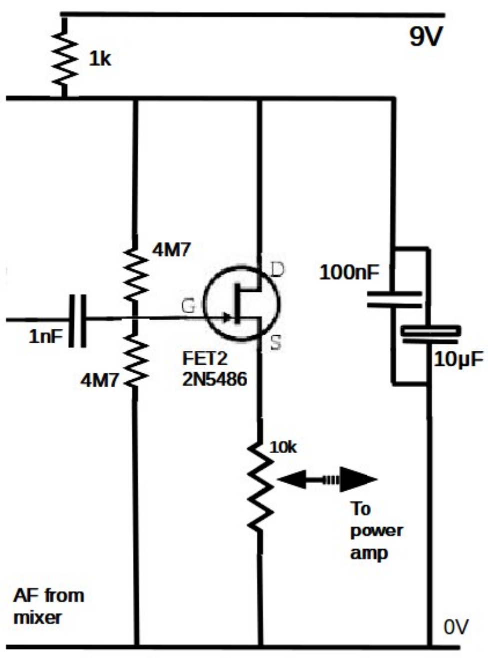

Fig 4.5.3.1: The pre-amplifier sub-module (this circuit will appear again as part of in the final figure of this section, Fig 4.5.3.3.5). There are two parallel decoupling capacitors because electrolytics, such as the 10µF one here, tend to have inductance significant enough to present high impedence to radio frequencies. On its own, the 100nF capacitor would be fine at high frequencies but would present too high an impedance to be efficient at decoupling low audio frequencies. Neither can do the required job of geneal decoupling on its own, but the combination of the two capacitors will perform the required task of grounding all signals on the supply line well enough.

4.5.3.2: The voltage-controlled attenuator sub-module.
The operation of the voltage-controlled attenuator is simple. The incoming audio signal (isolated from any upstream dc by the 1200nF capacitor) is placed across a potential divider consisting of one fixed resistor and one FET (Fig 4.5.3.2.1).

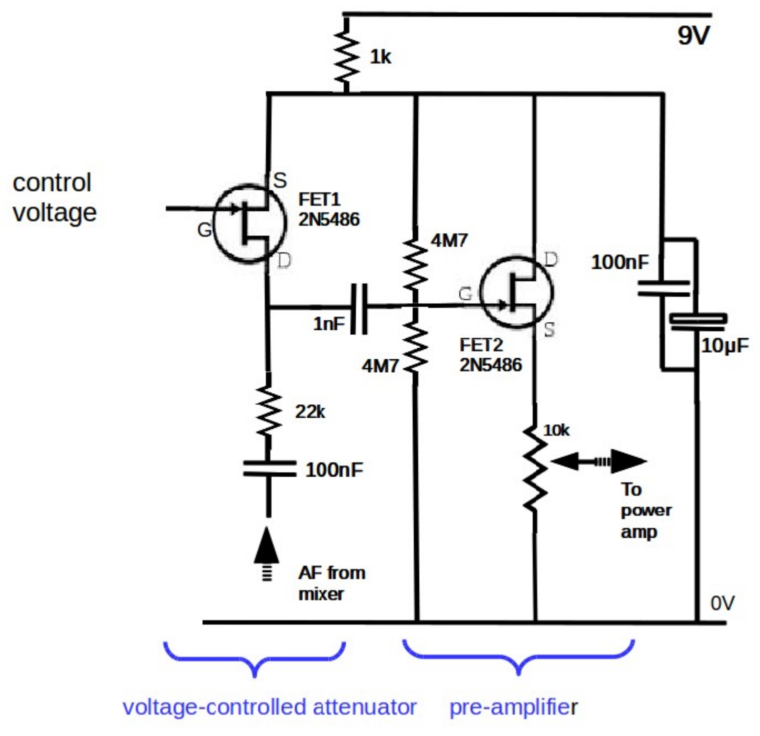

Fig 4.5.3.2.1: Using a FET in a potential divider to attentuate the AF signal (this circuit will appear again as part of in the final figure of this section, Fig 4.5.3.3.5). The choice of 22kΩ for the resistor in series with the incoming audio frequency (AF) signal is explained in the main text below. Please note the unusual 'way up' the FET is connected: having the source at the top is correct and not a misprint.

That this arrangement is a simple potential divider may not be obvious because the FET points 'up' to the positive supply not directly to ground. Remember, though, that from the point of view of a signal, the decoupling capacitors at the extreme right of the whole circuit diagram pass ac so well that make the dc positive supply rails indistinguishable from ground. 'Shorting' a signal up to the supply side is the same as shorting it to ground. FET1 is connected the way up it is ('upside-down') because of this: note that its source-drain path conducts only ac signals – there is no dc path out of the drain.

The effective resistance presented by the FET is controlled by the voltage between gate and source, as shown in Fig 4.5.3.2.2 below.

Fig 4.5.3.2.2: the relationship between conductivity of the FET and gate-source voltage, taken from the manufacturer's own data sheet (under fair use policy). Using the 15V figures as a guide, when the gate and source are at the same voltage, the device conducts well (dividing voltage by current gives a resistance of around 1kΩ). When the gate becomes negative compared to the source, the device conducts much less well, resistance rising to at least 20kΩ.

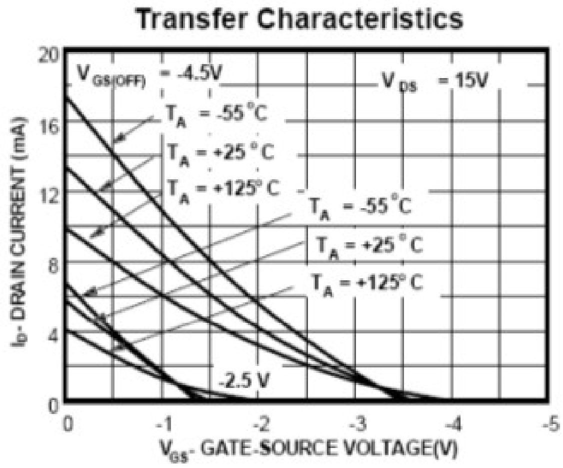

If we can arrange a a frequency-to-volatge converter that will provide a controlling voltage to the FET's gate of between 0V and about 4V below that of the Source terminal (ie below the supply line shared by all sections of this circuit), then we can expect the resistance of the FET to range from about 20k to 1k. This can be used as a guide to the value of the fixed resistor in the potential divider. Too low a fixed resistor, and the FET will be incapable of 'grounding' all of the current that can flow through the fixed resistor and attenuation will be trivial. Too high a value, and the FET will cause massive attenuation of the signal (taking it to 'signal ground') all the time. A sensible choice is a value close to the maximum resistance of the FET. Feeding the signal in via a 20k resistor will provide a 50% voltage attenuation at the point between resistor and FET when the FET has a resistance of 20k, and will provide a voltage attenuation of 95% when the FET has a resistance of only 1k. That's a 20:1 span of voltage attenuation, or a 400:1 span of power reduction (power goes with voltage squared), or expressed in decibels, that's a 26dB range. That is compatible with the range on Clara Rockmore's recordings. Note that this resistor value (22k, the nearest I had to 20k) is very different from those of the web design on which this build is based.

4.5.3.3: the frequency-to-voltage converter.

Proximity of a the player's left hand to the volume half-capacitor causes a shift in frequency of around 1% from no-hand to hand-right-down (see sections 4.5.4 and 4.5.5). To make best use of the voltage-controlled attenuator described in section 4.5.3.2, we need to build a device to translate this 1% swing in frequency to a 95% swing in dc voltage. The obvious analogue method for doing this is to apply the incoming signal to a resonant circuit. A classical example of a resonant circuit is a parallel combination of capacitor and inductor, in which energy is stored alternately in the electric field of the capacitor and the magnetic field of the inductor (Fig 4.5.3.3.1).

4.5.3.3.1: background to resonant filters.
Reader's navigational note: if you are completely familiar with resonance, f_r and Q, skip to 4.5.3.3.2

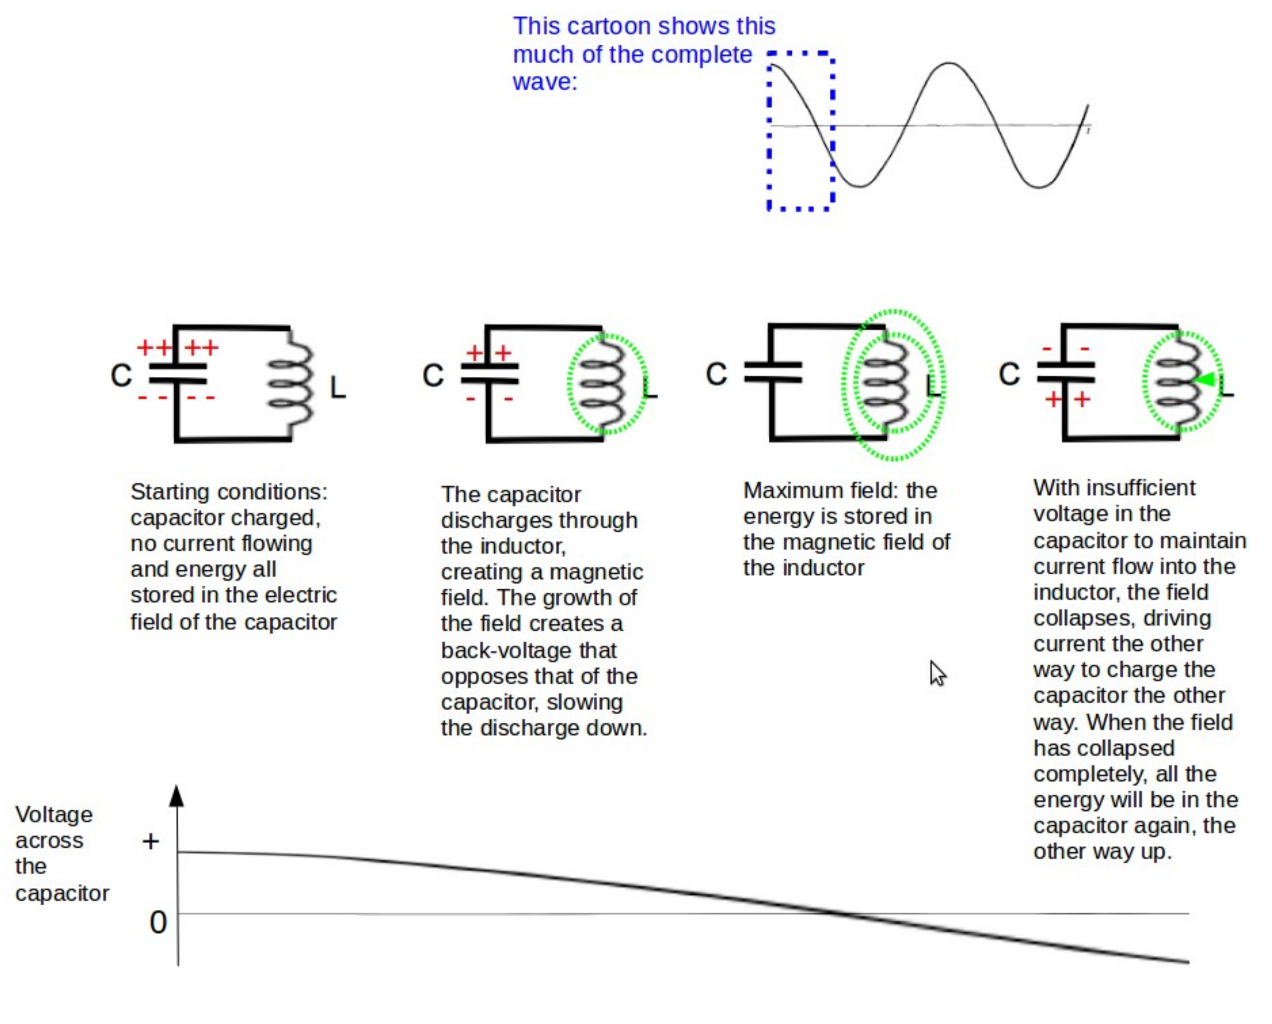

Fig 4.5.3.3.1: resonance in an LC parallel circuit.

A resonant circuit such as the one above is characterised by two values – f_r, the frequency at which it resonates best ('resonant frequency' / 'centre frequency'), and the quality factor, *Q*, which measures how fussy is is about frequencies being precisely at the centre frequency. The existence of a centre frequency can be understood by consdidering the degree to which each component of the L-C combination impedes the flow of energy into and out of it due to its reactance, and how this reactance depends on frequency.

The standard formula for reactance of an inductor of L Henrys inductance at a frequency of f Hertz is $Z = 2\pi fLj$ ohms, where *j* is the suffix indicating this is a reactance, not an ordinary resistance (mathematically minded readers will recognize that *j* actually denotes something more subtle then this but, for the purposes of this document, '*j* is the suffix indicating this is a reactance' will do). The formula for the impedance of a capacitor is $Z = -j/\ 2\pi fC$ ohms. The total reactance of the system is therefore $2\pi fLj\ \ -j/\ 2\pi fC$ ohms. At one particular frequency, reactance will be zero:

$$2\pi fLj\ \ -j/\ 2\pi fC = 0.$$

Rearranging this in terms of f (f_r as we will now call it) we get the well-known

$$\mathbf{f_r = 1/ 2\pi \sqrt{LC}}$$

The steepness with which the willingness of an LC circuit to oscillate falls off with mistuning of the input signal is described by the quality factor, Q. This is easiest to understand in the context of a graph showing what happens when a resonant circuit is driven at a range of frequencies (Fig 4.5.3.3.2):

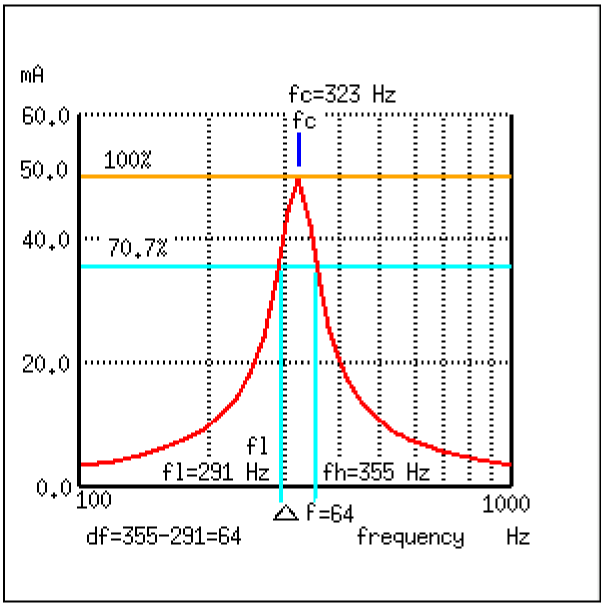

Fig 4.5.3.3.2: the relationship between Q and the steepness with which the willingess of a resonant system to accept frequencies falls away with mis-tuning. In this graph (copied from allaboutcircuits.com, under the terms of their Design Science licence), the y axis represents the amplitude of oscillations in a resonant system when fed with difference frequencies (x axis). The red line shows a typical shape for a simple LC system. The blue line shows the bandwidth (see main text), in this case 64Hz.

In this diagram, the centre frequency is labelled f_c instead of f_r: tghey mean the same thing.

In a simple system (eg LC), the shape of the red line is predictable but the width and steepness of the peak will vary with Q (as if the graph is stretched sideways and compressed up-down). It is therefore possible to capture the steepness of the graph in one simple measure, the bandwidth, which is defined as the range of frequencies over which the amplitude (volts or current, it does not matter) is greater than or equal to 70.7% of the amplitude at the centre frequency. Why the weird figure of 70.7%? Because the bandwidth is really defined as the range at which the amplitude of *power* is at least 50% of that at the centre frequency. Power is proportional to the square of voltage (or current) amplitides ($P = IV = I^2R = V2/R$ where R is the resistance in the system – Ohm's law), and the square root of 0.5 (=50%) is 0.707 (70.7%). That's where the peculiar number comes from.

The concept of bandwidhe can be useful in engineering communication systems to be narrow-band, broad-band etc, but it is not very helpful generically because a bandwidth of 3kHz on a centre would be much less selective, in relative terms, on a centre frequency of 10kHz than it would on a centre frequency of 100MHz. A more helpful generic term is one that brings bandwidth and centre frequency together, which is what Q, the quality factor, does:

$$\mathbf{Q = f_r / bandwidth.}$$

Resonant systems with different values of Q show quite different slopes, high-Q equating with a larger response at the centre frequency and a more fussy one in terms of the difference between a perfectly tuned and a near-miss frequency.

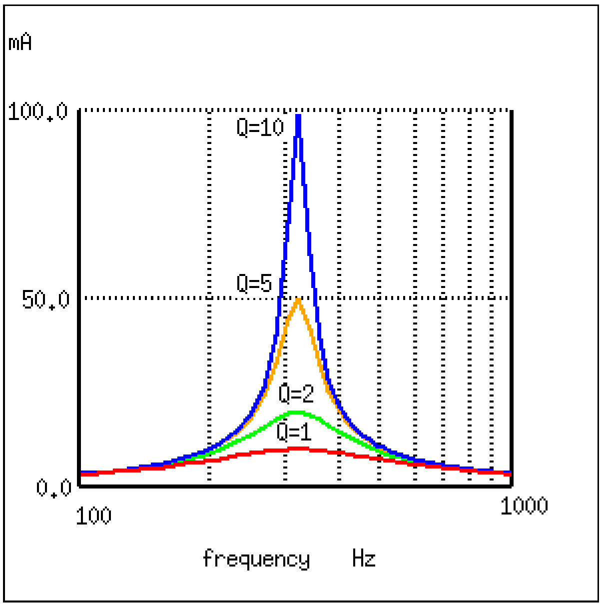

Fig 4.5.3.3.3: This graph (copied from allaboutcircuits.com, under the terms of their Design Science licence), shows the responses of resonant circuits (amplitude up the y axis, frequency across the x), all with the same centre frequency but different values of Q.

Obviously, to get a steep slope so that a small change in driving frequency will make a large difference in amplitude in the resonant system (which is what we need for the volume control system of the Termenvox), **we will need a resonant system with a high Q.**

In any resonant system (not just in electronics), the value of Q is predictable from the rate at which the system loses the energy that has entered it.

$$Q = 2\pi \times \frac{\text{Energy Stored}}{\text{Energy dissipated per cycle}} = 2\pi f_r \times \frac{\text{Energy Stored}}{\text{Power Loss}}.$$

Most textbook treatments present the idea of energy loss in terms of power dissipated through a resistor. A resonant circuit consisting of an inductor and capacitor and no deliberately added resistors will still have electrical resistance in the interconnecting wires and, more seriously, in the wires coiled up in the inductor. Ferrite inductor cores can also dissipate energy as magentic fields drine 'eddy currents' in the core. This internal resistance limits the maximum Q that can be achieved (and this limit is higher with high-quality components). Textbook treatments also introduce the idea of deliberately adding resistance when one wants to have a low Q and a deliberately wide bandwidth, as for a broad-band filter.

What texts often omit to say, taking it as obvious I suppose, is that power is also lost to the system to whatever devices are connected to it, for example as current moving into the base of a transistor connected to it. This is important, because it means that the real-world 'loaded' Q of a resonant system connected to an output falls short of the ideal Q of the same LC resonator in isolation.

4.5.3.3.2: Choosing a resonant system for the Termenvox.

In building the Termenvox from a design on the web, I blindly followed the suggestion for a circuit in which a single LC resonator was fed with the output of the volume variable oscillator, and the amplitude developed across the resonator was rectified into dc to drive the attenuator. The result was disappointing – the roughly 1% swing in driving frequency caused only about a 50% change in output voltage, not the 95% that is necessary for the dynamic range we require. The problem was probably a combination of resistance in the inductor, and loading from the following circuit.

Termenvox volume circuits based on the original RCA valve-driven models do use LC circuits, but in the context of oscillators in a 'grid-dip' confugiration in which the voltage on a valve grid swings strongly as the resonance of the volume proximity detector system changes. This is a smart idea, because it gives the varying voltage output directly from the oscillator. Had I thought of it (I am kicking myself for not doing so, as I frequently use grid-dip oscillators as test instruments), or had I studied the RCA design before reading around to try to debug the device I had already built from the web design, I may well have followed it (using a FET in place of a valve). Having already built the volume variable oscillator to give an output primarily in terms of frequency shift, I pressed on with the idea of increasing the 'Q' of the resonant system for turning this frequency shift into a large change in amplitude.

It is possible to make higher Q filters by using combinations of multiple capacitors and inductors, but modern consumer radio manufacture has driven the development of cheap and plentiful 3-terminal ceramic resonators with quite high Q and a centre frequency of 455kHz (a standard intermediate frequency in superheterodyne radios). The ceramic filter specified in Fig 4.5.3.1 is typical: it has a centre frequency of 455kHz, and a frequency response as in Fig 4.5.3.3.4. It can be seen that, if the volume variable oscillator is set to run at 454kHz under no-hands conditions, even a 0.5% reduction in its frequency (2.25kHz, taking the frequency down to 451.75kHz) will result in a reduction of output from 50% input to about 5%, a 90% fall. The full 1% change would result in a fall of around 90%. This is much better! Also, the resonator has input and output impedances of just 2kΩ, so loading was unlikely to be an issue. There is a risk that the Q for this filter will be too high, but this can be softened by adding resistance in the ground line of the filter. I found empirically that 3.3k was fine: it would be possible to make this resistance a variable one (rheostat), so that the musician can control how savagely the volume control system silences the sound.

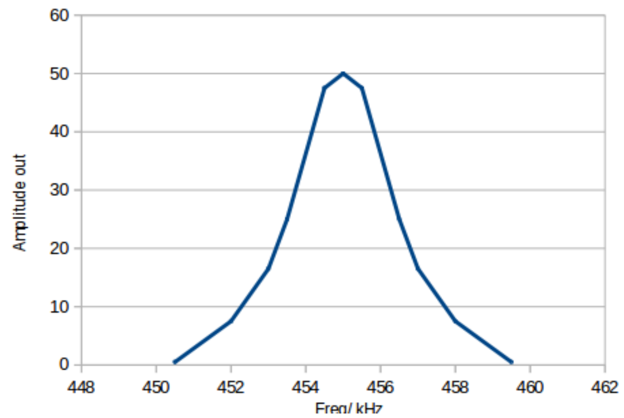

Fig 4.5.3.3.4: Output voltage of the ceramic resonator, as a proportion of input voltage, at a range of frequencies.

In principle, an isolated LC resonator can look as good, but not when it is loaded by a rectifier. The measurements left were made into a 2k load, which is a good match for the manufacturer's claimed output impedance of the resonator.

The frequency-to-voltage converter was therefore rebuilt to the form shown in Fig 4.5.3.1. The output from the volume variable oscilaltor is passed in to the input pin of the ceramic resonator. The output from the output pin is taken through a 2.2k resistor (to respect the impedance requirements), and a blocking capacitor (to isolate from dc voltages), and then passed directly to two diodes. D2 shorts negative halves of the cycle to ground, and D1 passes positive halves to a reservoir capacitor, which will be drained steadily by the 15k resistor across it. I used small signal diodes. Bipolar transistor Q1 (bipolar is fine here because the output impedance of the ceramic resonator is so low) amplifes the dc voltage at its base. It is connected in common-emitter configuration, giving voltage gain and inversion. When there is no signal coming out of the resonator, Q1 is 'off' and the voltage at point A in the diagram is up at the supply voltage at the top of the collector resistor. The gate of FET1, connected to this point by a resistor (and a capacitor to shunt any remaining ripple to ground), is therefore not pulled negative with respect to the source, and FET1 has low resistance, shorting the audio signal to ground. As more signal comes from the ceramic resonator (the player withdraws her hand from the volume half-capacitor, allowing the volume variable oscillator to return towards 454kHz), the voltage on Q1's base rises, admitting current, and Q1 beings to conduct. The voltage at point A falls with respect to ground or, as the FET1 sees things, the voltage on its gate becomes negative with respect to its source. FET1 therefore develops more resistance and shunts less audio signal to 'signal ground', and more audio signal is left to reach the pre-amp, power amp and speaker.

The whole circuit of the frequency-controlled amplifier module is in Fig 4.5.3.3.5:

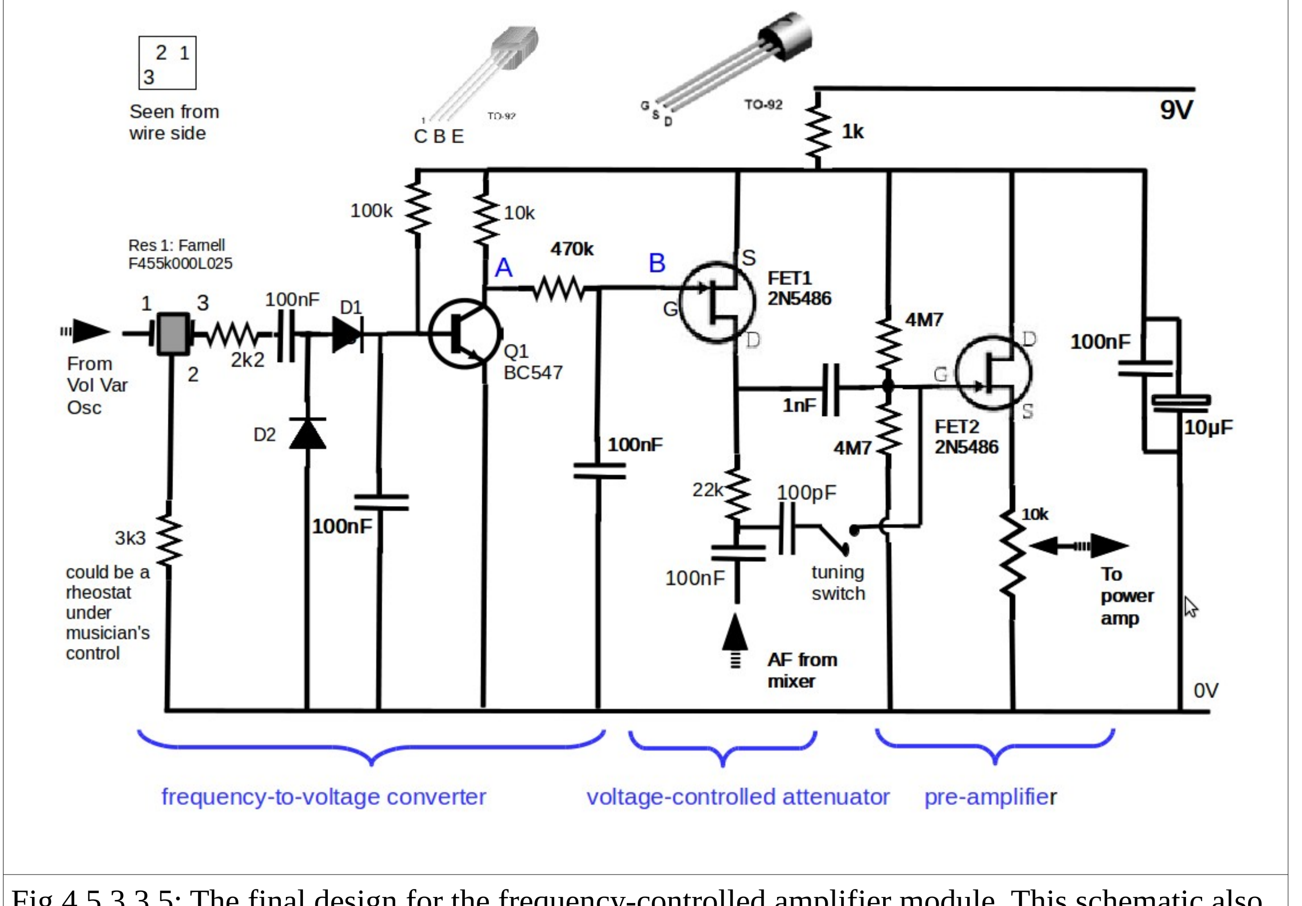

Fig 4.5.3.3.5: The final design for the frequency-controlled amplifier module. This schematic also shows the position of the tuning switch, an optional extra that makes life easier both for the contstructor and the musician. The swtich is closed for 'tune' and open for 'play'.

Testing (of the whole VCA module): Before connecting Q1 to point A, connect an AF signal generator to the 'from mixer' point, and check that you hear the signal (this assumes you have built and connected the power-amp and speaker). If you do, use a temporary resistor or variable voltage source to pull point A down to ground: by the time you reach about 1.2V from ground (about -4.2V from the FET1 source terminal), the speaker should be completely silent. Assuming that's OK, connect the crystal and diodes. With no RF signal, the speaker should still be silent. With an external RF signal of 454.0kHz (2Vp-p will do) applied to the crystal input, there should be significant AF reaching the speaker. Pulling the RF signal down to 451.5kHz should progressively reduce the volume of the AF signal reaching the speaker to nearly zero.

4.5.4 The volume oscillator

The volume oscillator provides the RF input to the resonant filter of the VCA module. When the detector module (described in section 4.5.5 below) is not signalling the presence of a nearby hand, the volume oscillator needs to provide a steady 454kHz . This frequency is then pulled down by the proximity of a hand to the detector module, by an action that will be described in section 4.5.5.

Typical RF oscillators have their frequencies set by resonant L-C circuits such as that in Fig 4.5.3.3.1. Unlike filters, oscillators have to generate, rather than simply respond to, oscillations. Energy lost in and from the L-C resonator has to be replenished, and this is done by sampling the signal from some place 'part-way-up' the resonant circuit, amplifying it, and feeding it in augmented form back into the top. In order for the sampling to work, there are two standard ways of re-arranging the simple L-C circuit. In one, invernted by Ralph Hartley in 1915, the inductance is split between two separate inductors in series and the sample is taken from a point between them. In the other, invented by Edwin Colpitts in 1918, the capacitance is split between two separate capacitors and the sample is taken between them (Fig 4.5.4.1).

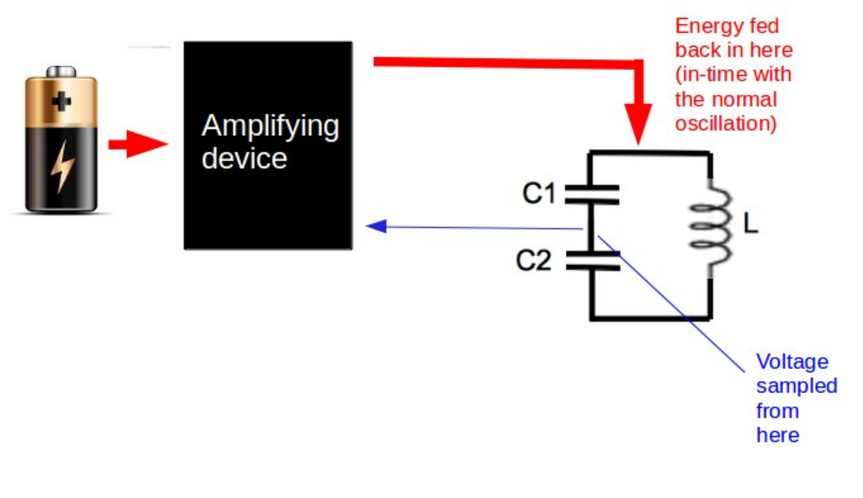

Fig 4.5.4.1: Colpitts' idea of splitting the capacitance of an L-C resonant circuit so that a sample of the oscillating signal can be taken, amplified and fed back into the system.

Since capacitors are smaller, cheaper and more easily available than inductors, it makes sense to use the Colpitts form for the Termenvox oscillators, including the volume one being discussed here. There are three important points to be made about the “Colpitts capacitors” C1 and C2. The first is that their values are much higher than the single capacitor needed for a simple L-C resonator, because the capacitance of two capacitors in series is calculated by

$$C_{tot} = \frac{1}{\frac{1}{C1} + \frac{1}{C2}} = \frac{C1C2}{C1+C2}$$

The second point is that the proportion of energy taken to be amplified, usually expressed as the feedback ratio C1/C2, depends on the relative size of the two capacticances. Where robust and reliable oscillation is required of a bipolar transistor and the risk of harmonic generation does not matter (as it does not here), a feedback ratio of about 8% is desirable.

For a fixed-frequency oscillator and a known value of inductance, these two points are in principle enough to allow the values of C1 and C2 to be calculated. Truly fixed and unadjustable oscillators are, however, rare and it is common for some adjustable capacitance to be included, in the form of a trimmer, so that the frequency can be set to exactly that required. Making C1 or C2 adjustable, or placing a trimmer in parallel with either one of them, would alter both the total capacitance and the feedback ratio, which would be a bad idea (it could, for example, result in feedback becoming too low to sustain oscillation). For this reason, adjustable capacitance is generally placed directly across the inductor, creating the arrangement in Fig 4.5.4.2:

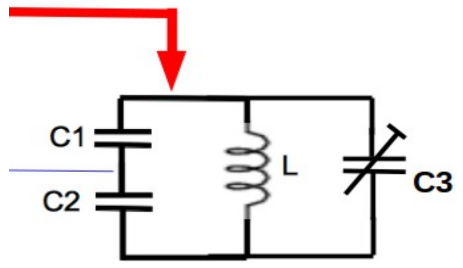

Fig 4.5.4.2: adding an adjustable capacitance, C3, to tune a Colpitts network. Total capacitance is now

$$C = C3 + \frac{C1C2}{C1 + C2}$$

Resonant frequency is still $f = \frac{1}{2\pi\sqrt{(LC)}}$

(There is also a fourth point, about phasing and C2, but this is dealt with in a few paragraphs' time).

For the Colpitts network of the volume oscillator, I wanted to have lots of 'wiggle-room' for experimentation with the value of C3, so I split the total capacitance roughly equally between C3, and the 'Colpitts capacitors' C1 and C2. With an aim of tunability around a mid-point of 454kHz, and an inductor of marked 330µH that actually measured 265µH on a meter, I decided on C3 being 200pF in parallel with a 60pF trimmer, C1 being 285pF and C2 being 3.59nF. These values yield a feedback ratio of 7.7%, a “C1,C2” capacitance of 265pF, a total capacitance from 479pF to 519pF and resonant frequency in the absence of any other influences between 426 and 453kHz. My own experience is that common-base Colpitts oscillators always seem to run about 5% higher than calculations suggest (see below), so a calculated range of 426-453 kHz should yield a real range of about 447-475kHz.

Having established the basic parameters of the resonant network, the next question is the arrangement of the amplifier. Field effect transistors (FETs) are a popular choice for oscillators but most are vulnerable to damage by static electicity. There will be a direct connection to this circuit from the metal volume half-capacitor of the volume detector, and this could easily be exposed to static electrity (for example, because the player has just taken off a pullover and then touches the half-capactitor). For this reason, the more robust bipolar type of transistor seems appropriate.

Transistors can amplify in three modes;

- The common emitter mode, in which the emitter is typically tied to ground, the signal is applied between ground and base, and the load appears between collector and supply, has the following features: medium-value input impedance, high voltage gain, and inverted output and an output impedence approximately equal to the load resistor. The common emitter mode is the one almost always drawn in elementary introductions to electronics.
- The common collector ('emitter follower') mode, in which the collector is tied to the supply,

the signal is applied between base and ground and the load appears between collector and ground, has the following features: very high input impedance, unity voltage gain, non-inverted output and very low output impedance.

- The common base mode, in which the base is grounded, as far as signals are concerned, by a large capacitor, an emitter resistor is used to raise the voltage of the emitter, the input is applied across the emitter resistor and the load appears between collector and supply, has the following characteristics: very low input impedance, strong voltage gain, unity current gain, a non-inverted output and output impedence approximately equal to the load resistor.

Of these, the most useful for the oscillator is probably the common base: the phase inversion of the output in a common emitter design would be a complication, and the very high input impedance of a common collector design would risk high impedence signals being picked up from the half-capacitor of the hand proximity detector (which could act as a a receiving antenna), and amplified. The common base has a very low input impedance that would squash these small signals. Also, the common base arrangement does not suffer the Miller effect (a multiplication of a transistor's natural small base-collector capacitance by stage gain).

The basic Colpitts design applied to a common base amplifer would therefore look like this:

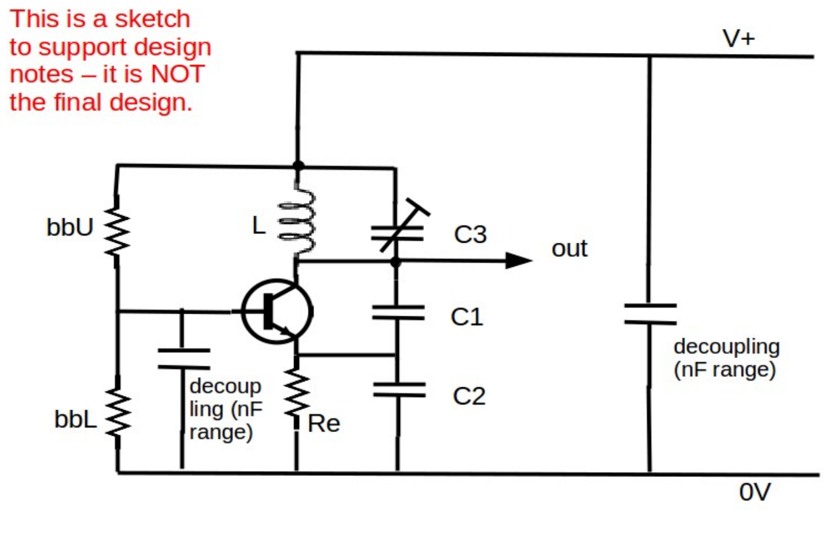

Fig 4.5.4.3: the basic design of a Colpitts oscillator using a transistor in common base mode. C3 is really a trimmer in parallel with a fixed capacitor, as explained in the text.

The Colpitts network may look different because the inductor is also being used as a collector load, but it is the same as far as an RF signal 'sees' it: the decoupling capacitor across the supply lines 'looks like' a dead short to RF, so the top of the inductor 'feels' connected to the 0V line, across the Colpitts capacitors C1 + C2. Similarly, the base is tied to 0V, as far as RF is concerned, by the base decoupling capacitor. The base-emitter voltage is therefore that developed across the emitter resistor Re. Base biasing resistors bbU and bbL clamp the dc value of the base bias in the usual way.

For this basic circuit to be of practical use, some additions are needed. The first is that the 'V+' needs to be modest, to keep heat generation as low as possible. The standard line voltage of 9V used as a power bus is too high, so a 'dropper' resistor is needed in series with the whole module.

The second change is to provide a buffer between the output of this module, and the input of the VCA module, so that changes of impedence in the VCA do not feed back to and affect the frequency of the volume oscillator. There will be no need for voltage amplification, just the transformation of a medium impedence into a very low one. A second transistor in common collector ('emitter follower') mode is the obvious choice.

The third addition is to allow the player to make small adjustments to the natural frequency of the oscillator, in order to compensate for different playing environments that will alter the 'pull' of the detector (see section 4.5.5). Direct access C3, by C3 being a radio tuning variable capacitor, for example, would be too 'sensitive' (in the sense that tiny movements would move the frequency too much, making proper setting frustrating). Something more subtle is needed. One 'extra' capacitance that is present in any transistor circuit is a small capacitance between collector and base. Given that the base is grounded (to RF) in the circuit, this capacitance appears in parallel with {C1+C2} and C3. According to the datasheet for the BC547 transistor to be used for this module, the value varies in a voltage-depedent manner, being about 5pF when the collector-base voltage is zero, about 1.2pF when the collector-base voltage is 9V. Usefully, in a common base circuit the "Miller effect", which would multiply the value of this capactitance by gain an a common emitter circuit, is absent so we can use the raw data sheet figures. Providing a means for the user to alter the voltage on the base of the transistor would therefore provide a means to make a small alteration of the total capacitance of the network. This can be done with a potentiometer (mounted inside the mint tin, with a long plastic shaft connecting to the control knob on the front of the instrument). The circuit carrying these modifications appears below as Fig 4.5.4.4:

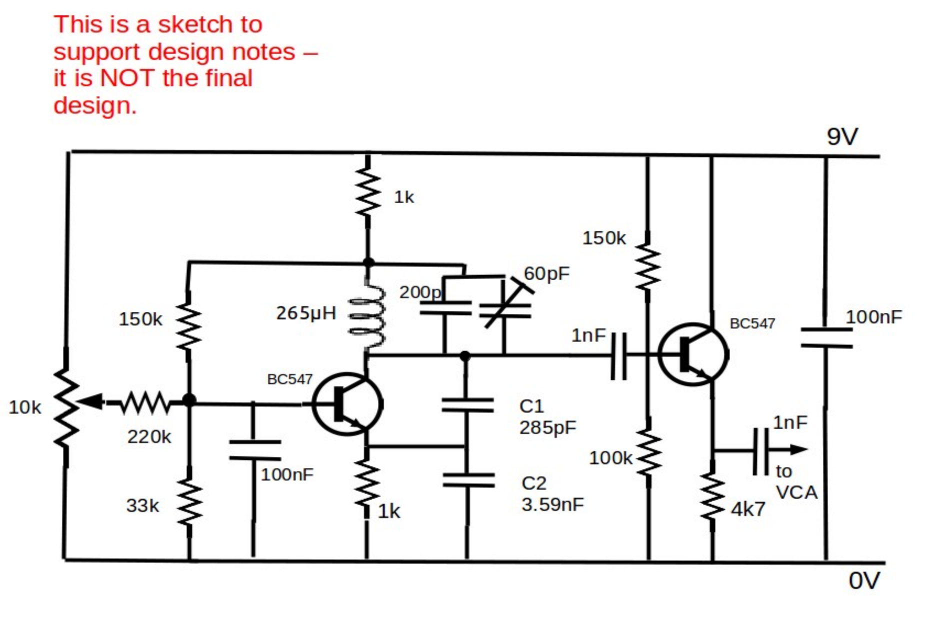

Fig 4.5.4.4: The Colpitts oscillator with the addition of a supply coltage dropping resistor, voltage-controlled tuning, and a common-collector buffer amplifier. C1, C2 and C3 need to be NPO/ COG types for thermal stability: I had a limited range of these types, so made C1 of a 270pF and a 15pF capacitor in parallel, C2 of a 2.2nF, a 1nF and a 390pF capactitor in parallel, and C3 of 2x 100pF in parallel (but this last value was altered anyway for the final version – see Fig 4.5.4.6).

Building the circuit above, and testing it first at the early morning temperature of my work area (0.5°C, the usual winter temperature of my house before I have lit the fire) and then in a room warmed to 20°C, revealed a serious flaw: the frequency altered with temperature, and did so really seriously – by about 12kHz. Aiming a haridryer at it could move the frequency by up to 20kHz. In something that needs to be stable to a few hundred Hz, this temperature dependence was clearly unacceptable. To investigate the problem, I took twins of the capacitors and inductor (I tend to but components in multpacks), connected them to an L/C meter, and measured their values in room temperature air and with a hairdryer pointing at them at close range (but not pointing anywhere near the L/C meter). The inductor, which was of a ferrite cored type, changed value by 2% between being at room temperature (5°C at the time) and being at a temperature I could touch but not for long (about 60°C). The capacitors, which were orndinary ceramic disc capacitors, changed by 19% over the same temperature range. Clearly they were at least a major part of the problem. I replaced C1,C2 and C3 them with NPO types (COG would show similar temperature invariance). The other capacitors are only for decoupling and their values are not critical.

This swap to NPO capacitors improved stability greatly, but it was still imperfect, the frequency changing from 457.4kHz to 456.2kHz over the first 10 minutes of operation.For the pitch system, this range would probably be tolerable (though not ideal) as musicians are used to having to re-tune for pitch drift in traditional acousitc instruments. Musicians are not, though, used to having to tune for volume, so more stability of the volume oscillator would be needed.

Careful observation of the changing frequency of the module with time, and also the voltage on downstream of the dropper resistor, revealed a suspicious relationship:

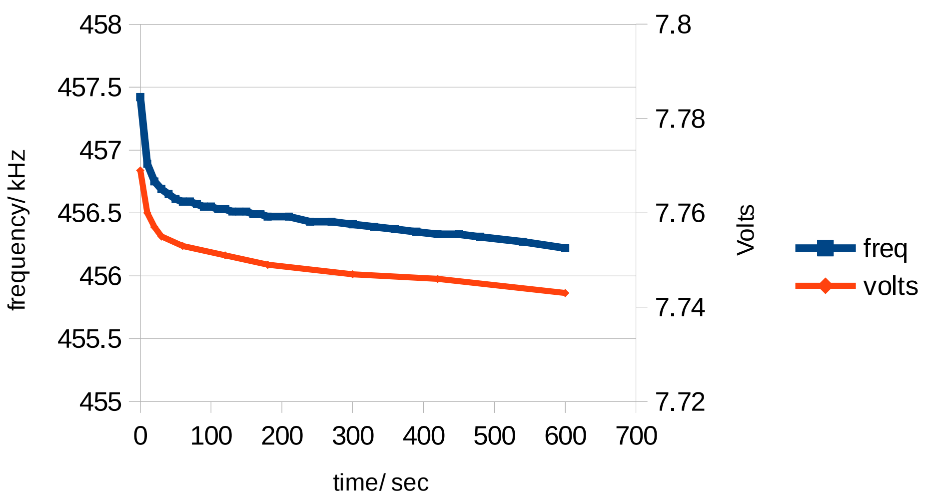

Fig 4.5.4.5: frequency and voltage at the top of the inductor correlate.

When two variables correlate, it is not formally clear which causes the other, or even whether both reflect a third process. That a change in voltage should alter frequency is, a priori, reasonable (because it will alter the difference between base and collector voltage and so alter the transistor capacitance – that same capacitance we are using for our potentiomoter-driven tuning system). Temporarily holding a 30pF capacitor across the inductor to reduce the frequency daramtically had no effect on voltage, allowing us to reject the reverse idea that changing frequency drives change in voltage. The cause of the changing voltage is presumably a slow warm-up of the transistor slightly increasing its current draw, and therefore elevating the voltage drop across the 1k dropper resistor. The simplest way to stabilize this is to place a Zener diode between the bottom end of the dropper resistor and ground; a 7.5V Zener makes very little difference to the voltage (which was naturally around 7.7) but does clamp it.

The final circuit is therefore as in Fig 4.5.4.6.

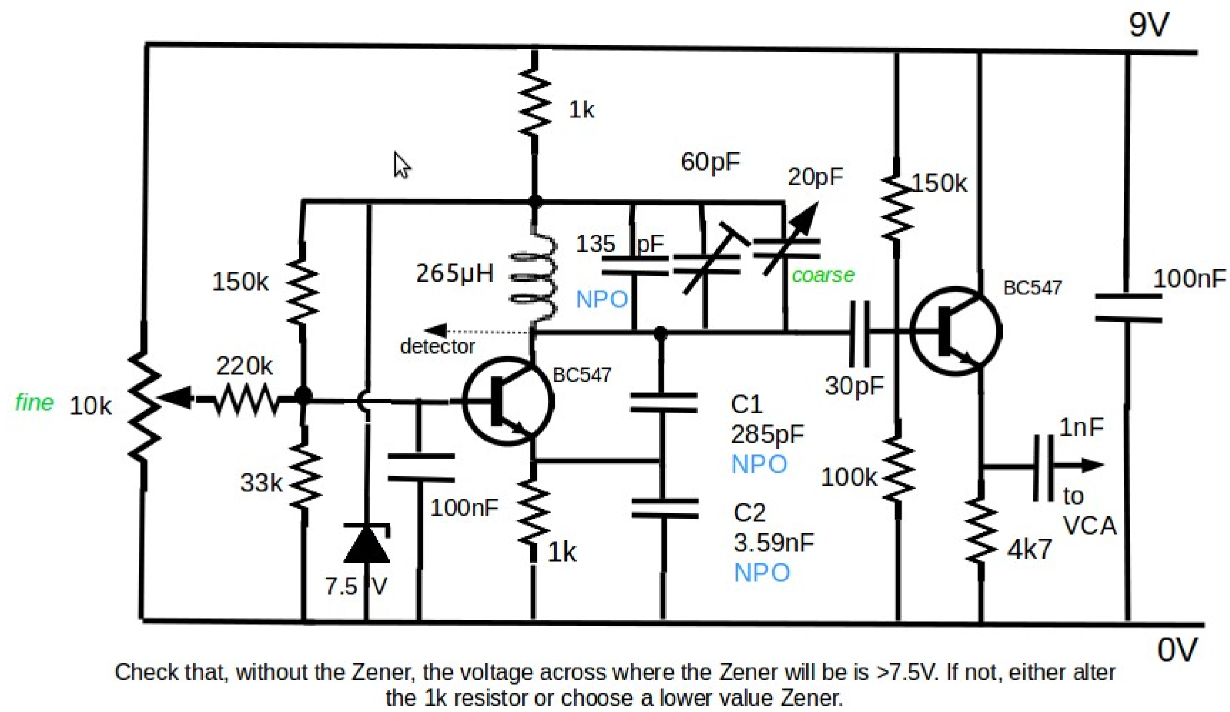

Fig 4.5.4.6: the final design for the volume oscillator. C1, C2 and C3 need to be NP0/ C0G types: as the exact values mentioned are not availbale in NP0/C0G – at least I could not find a supplier – C1 is actually made of a 270pF and a 15pF capacitor in parallel, C2 of a 2.2nF, a 1nF and a 390pF capacitor in parallel, and C3 of 2x 100pF in parallel, and the 135pF capacitor was made from 2x o 270pF in series.

Here is the frequency stability of this design (green) compared to the no-Zener previous design (blue):

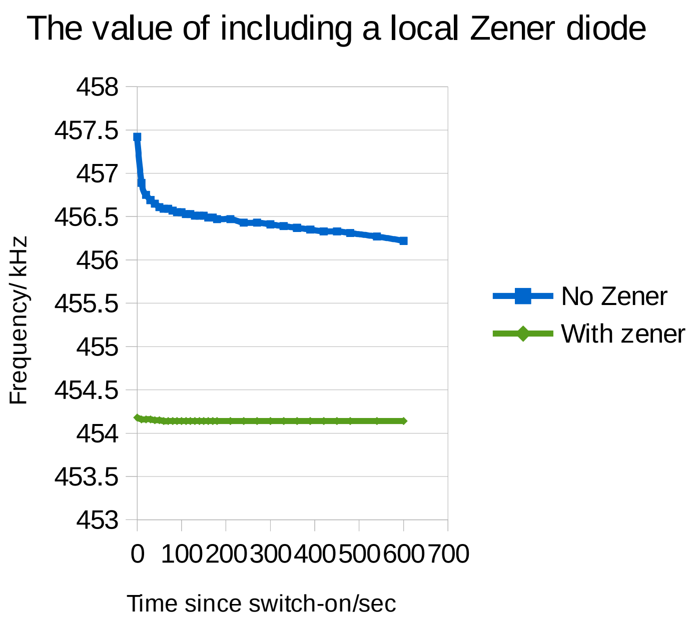

Fig 4.5.4.6: much-improved frequency stability resulting from the use of NPO capacitors and Zener stabilization of the local power supply to the oscillator.

It will be noticed that one more component has been added, a 20pF variable capacitor controllable from the front panel. This was added after the instrument was “finished” and was being used. It turned out that, while temperature stability was now OK, changes in relative humidity could cause the volume loop to pull the volume oscillator too far for the fine control potentiometer to re-tune it. The 20pF capacitor was therefore added as a coarse control. The 20pF capacitor chosen was the FM tuning section of a polyvaricon, with 2 x 20pF sections and 2 x 266pF sections, with additional trimmmers, from ebay (http://stores.ebay.co.uk/usefulcomponents?_trksid=p2047675.l2563). Any similar capacitor would do, but details of the one actually used is illustrated in Fig 4.5.4.7.

To set up the volume control oscillator, set the coarse and fine controls to their middle position and use the trimmer to check set it runs somewhere around 454kHz *without the proximity detector unit connected*. Don't worry about setting the frequency precisely yet, though, because the volume proximity detector will 'pull' it down a little even when no hand is nearby.

By the way, measure the frequency at the output of the buffer amplifier, not in the volume oscillator itself (frequency counter leads and input stages add capacitance – mine add about 15pF – so will alter the frequency of the oscillator to which they are connected).

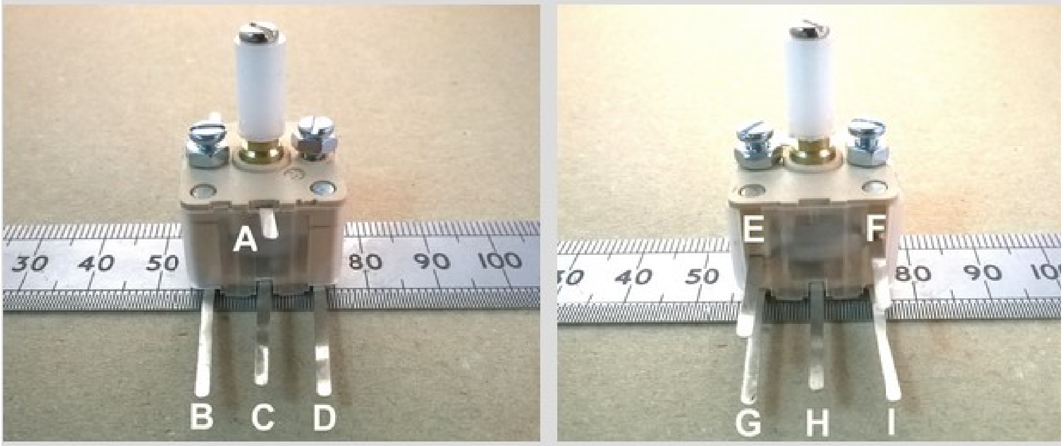

Connection Diagram Key

```
A  Common (GND)
B  VC1, 20pF + C1 15pF trimmer, FM
C  Common (GND)
D  VC2, 20pF + C2 15pF trimmer, FM
E  VC4, 266pF, AM
F  VC3, 266pF, AM
G  C4, 15pF trimmer
H  Common (GND)
I  C3, 15pF trimmer
```

Fig 4.5.4.7: pinout diagram of the polyvaricon. The 20pF capacitor used was the one between pins B and C: the 15pF trimmer in parallel with the variable capacitor was backed off to its minimum value (about 7pF total) before use, with the aid of a capacitance meter (there is no stop on the trimmer, so a meter is necessary). All unused leads were cut away very close to the body. This photo has been reproduced by courtesy of the supplier, Useful Components. *Important*: the shaft is connected to the 'ground' pins (which might be a problem if you use a metal box, because neither side of the capacitor connects to actual ground in the circuit. The collar at the bottom of the shaft can all-too-easily contact the inside of a hole through the case's front panel).

4.5.5 The volume proximity detector module.

The proximity detector module is 'really' just an addition to the resonant network of the Volume Oscillator but its action is complex enough that I find it easier to think of it as a separate control unit.

The 'player end' of the proximity detector is a metal 'half-capacitor' that projects from the side of the box. The player's hand forms the other half. The projecting metal gains capacitance from two sources. One is just the capacitance of any conductor against space itself (with no need for another nearby conductor), sometimes called 'self-capacitance'. The capacitance of a sphere of radius r, for example, is $4\pi\varepsilon_0 r$ and that of a disc is $8\varepsilon_0 r$, where ε_0 is the 'electric constant', $8.854\text{x}10^{-12}$ F/m (see E&M p109). When there is no hand near the half-capacitor, that is all the capacitance there is, neglecting for a moment the unavoidable capacitance of the socket into which the half-capacitor is plugged. When the hand comes close, it acts like another capactive plate, and the capacitance of this is $\varepsilon_0 A/d$, where d is the distance between the hand 'plate' and the half capacitor 'plate', and A is the area of each. For the dimensions of the volume half-capacitor used here (details in a moment), and using the statement that the self-capacitance of an annulus is $4\pi\varepsilon_0 r$ that appears in E&M page 114, self-capacitance can be estimated to be of the order of 10pF, and hand capacitance would add about 1pF at a distance of 1cm. These rough calculations can provide a foundation for design of the proximity detector, in the knowledge that they are rough and some fine adjustment can be expected to be needed at the end, based on empirical testing.

It is noticeable that, in the Термéн design, the half capacitor for the volume control is a loop rather than a rod. I think the reasons for this come down to geometry, and are helpful in giving a reasonably linear relationship between capacitance the the distance to the hand. Using the equations above, we can compare the shape of curve that results from a small object (eg a finger) being moved towards the centre of a ring-shaped half-capacitor (distance from finger to conductor can be calculated by Pythagoras' theorem), or towards the mid-point of a long, cylindrical one (this calculation requires integrating capacitances from short segments δl of the rod, along the rod). The results show curves that behave differently, the loop giving a more linear response when the finger is near, allowing delicate control of volume.

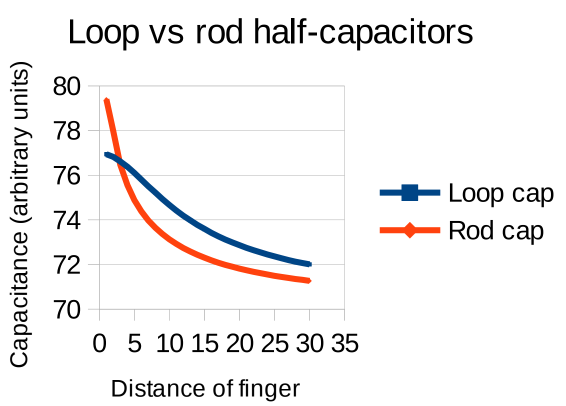

Fig 4.5.5.1: a rough estimation of the relationship between total capacitance on the half-capacitor (self-capacitance plus hand-plate capacitance) and distance between the centre of the half-capacitor and the hand.

There is a frequent misunderstanding, in descriptions of the Termen vox, that the player's hand controls the frequency of the volume oscillator (and the pitch one) by adding capacitance directly to the LC resonant circuit of the oscillator, with no additional components being needed. A moment's calculation will show why this will not work. The total capacitance of the Colpitts network is of the order of 500pF. The presence of the hand will add a 1pF capacitance even as close as 1cm, which would shift the frequency by only about 500Hz; at 10cm the shift would be only about 50Hz. In other words, taking this approach would affect the frequency too little, and too close, so that hands would be practically touching the half-capacitors to produce significant effects. This is exactly what happens in many 'toy' Theremin designs, but it is useless for a musical instrument and misses out on one of the cleverest things that Терме́н did.

The half-capacitor is not connected directly into the oscillator's resonant network but is instead connected via a very large inductance, effectively forming a series resonant circuit (Fig 4.5.5.2). This series resonant circuit intentionally breaks the rules of normal good design. In normal circumstances, designers take care to balance the risks to stable, predictable behaviour from inenentional stray inductances and capacitances in a circuit, the latter being more pervasive. In practice, this means that, as a rule of thumb, the value of capacitance in picofarads should be the same order of magnitued as the value of inductance in micro-Henrys. The Colpitts network in the Volume Oscillator follows this normal rule. The series resonant circuit of the volume proximity detector very definitely breaks these rules. It has very little capacitance (about 10pF), so has to have great capacitive reactance (= $-j$ / 2πfC). This means that, if the tuning of this network is changed by addition of modest capacitance, it can 'pull' the main Colpitts network very strongly downwards.

For the loop, I used an old wire loop antenna from a discarded portable television, which was made of stiff steel wire. I bent the very ends of this wire and fitted them with banana plugs, which fit into sockets on the side of the instrument. There are two sockets to hold the loop flat; only one is connected on the inside.Calculation suggests an inductance <u>of the order of</u> 10mH, but 'hidden' capacitances in the system, for example in the banana plugs and between the nearest parts of the loop to the case, will reduce capacitive reactance (= $-j$ / 2πfC, so falls as C rises), and less inductive reatcance (= 2πfLj) will be required to balance it for the same f (f = a little less than 454kHz).

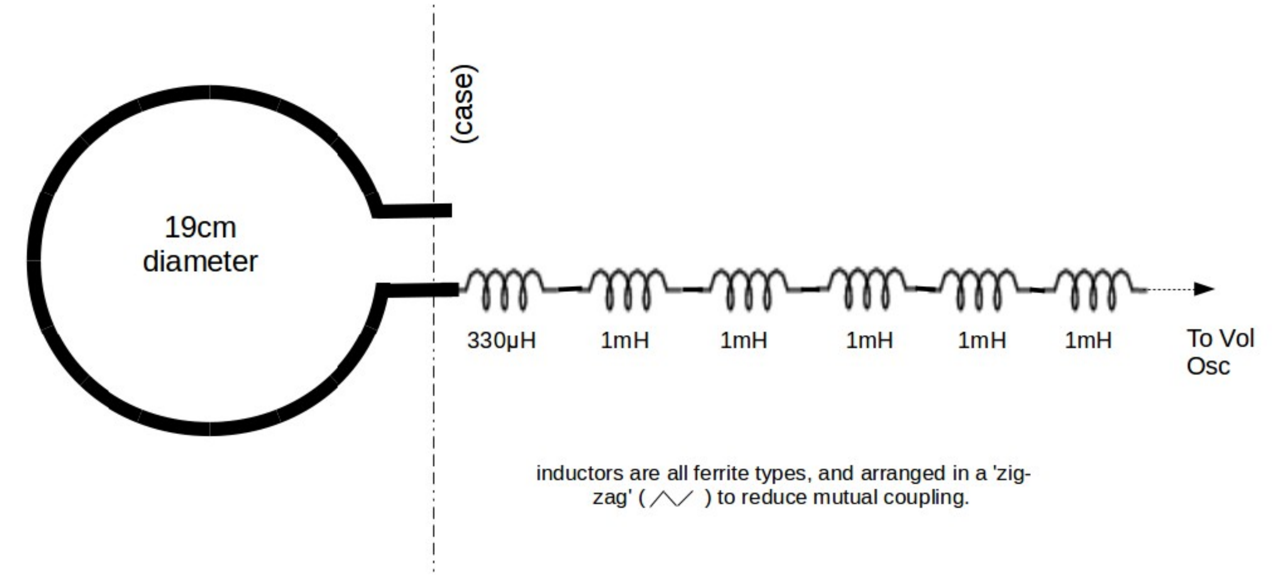

Fig 4.5.5.2: The volume proximity detector unit (different layouts, loops, choices of case etc will require different values of inductance). Electrically, the loop is a small capacitance to ground.

The procedure is to try adding inductance (inductors in series) until the series resonant circuit (loop plugged in but hands and other things well away from it) pulls the volume control oscillator's frequency down (not up) when it is connected, and shows a shift of at least 3kHz when a hand is moved to be close to the loop. Then use the trimmer capacitor in the volume oscillator to being the frequency of the whole system back to 454kHz, and check that the hand still pulls the frequency down to 451kHz or below. This process has to be empirical. In the case of my build, I needed 5.3mH of inductance, formed of five 1mH inductors and one 330µH inductor in series. Ensure that the hook-up lead is reasonably well located, so that it does not move about when the instrument is moved from place to place, and alter the inductive/ capactive connections between it and its surroundings.

4.5.6 The pitch reference oscillator

The pitch reference oscillator is designed to run at about 190kHz, slightly adjustable: it provides the reference note against which the hand-tuned pitch variable oscillator is heterodyned.

The pitch reference oscillator design follows closely that of the volume oscillator, so needs to particular additional explanation. The values of L, C1 and C2 are altered to produce the lower frequency of 190kHz. The exact frequency chosen does not matter much, but be careful that 454kHz is not a harmonic (an integer multiple) of it, to avoid the risk of audio-frequency heterodynes between the pitch and the volume oscillators, and to avoid the risk of harmonics of the pitch system feeding into the volume control. Obviously it is important that the frequency of the pitch reference oscillator and the pitch variable oscillator are the same.

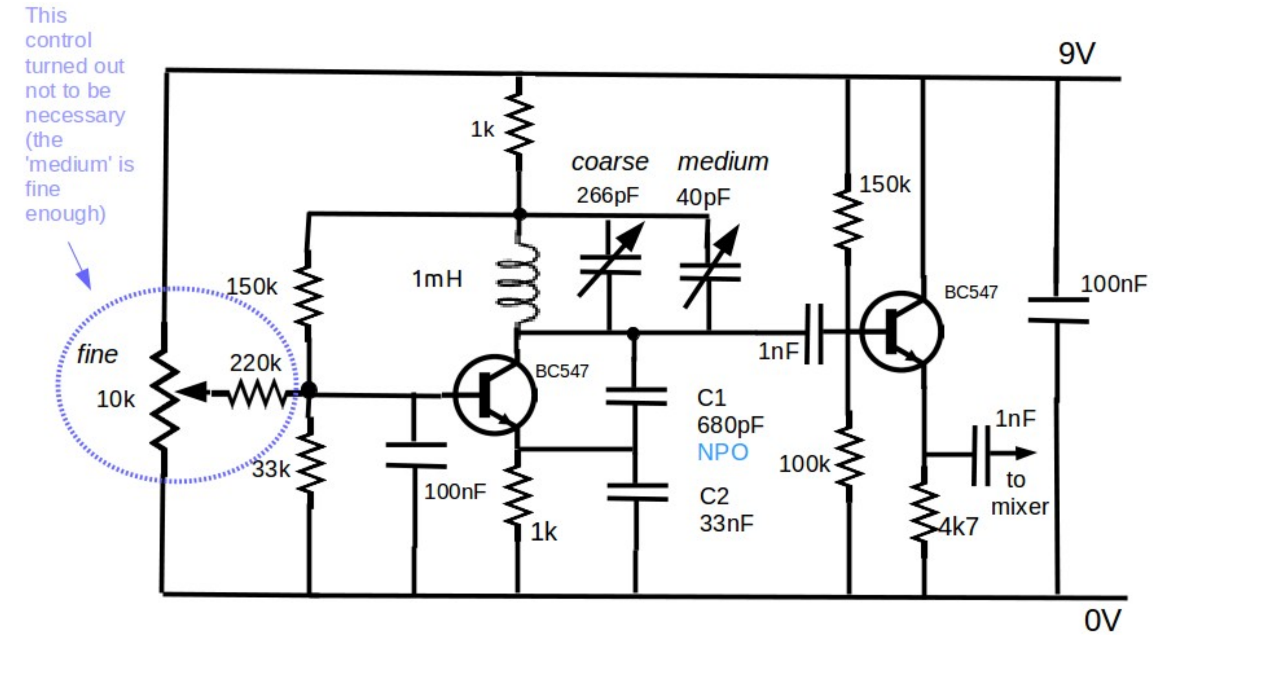

Fig 4.5.6.1: The pitch reference oscillator.

Frequency adjustment is made by three controls (although testing showed that only two were necessary). The first is very coarse – a 266pF section of a polyvaricon – and just sets the 'ball-park' frequency (the musician moves it until some kind of sound is heard). The second is a medium control, which is used to adjust the frequency to inaudibly low when no hand is near the pitch half-capacitor. It consists of 2 x 20pF sections of a different polyvaricon, connected in parallel. The third, which turned out not to be necessary, is a fine control achieved by using a potentiometer to alter transistor capacitance as in the volume control oscillator. If I build a second Termenvox, I will miss this control out: the medium control is quite fine enough.

4.5.7 The pitch variable oscillator

The pitch variable oscillator is based closely on the pitch reference oscillator, but without the potentiometer or variable capacitor to alter tuning. Set it to 190KHz.

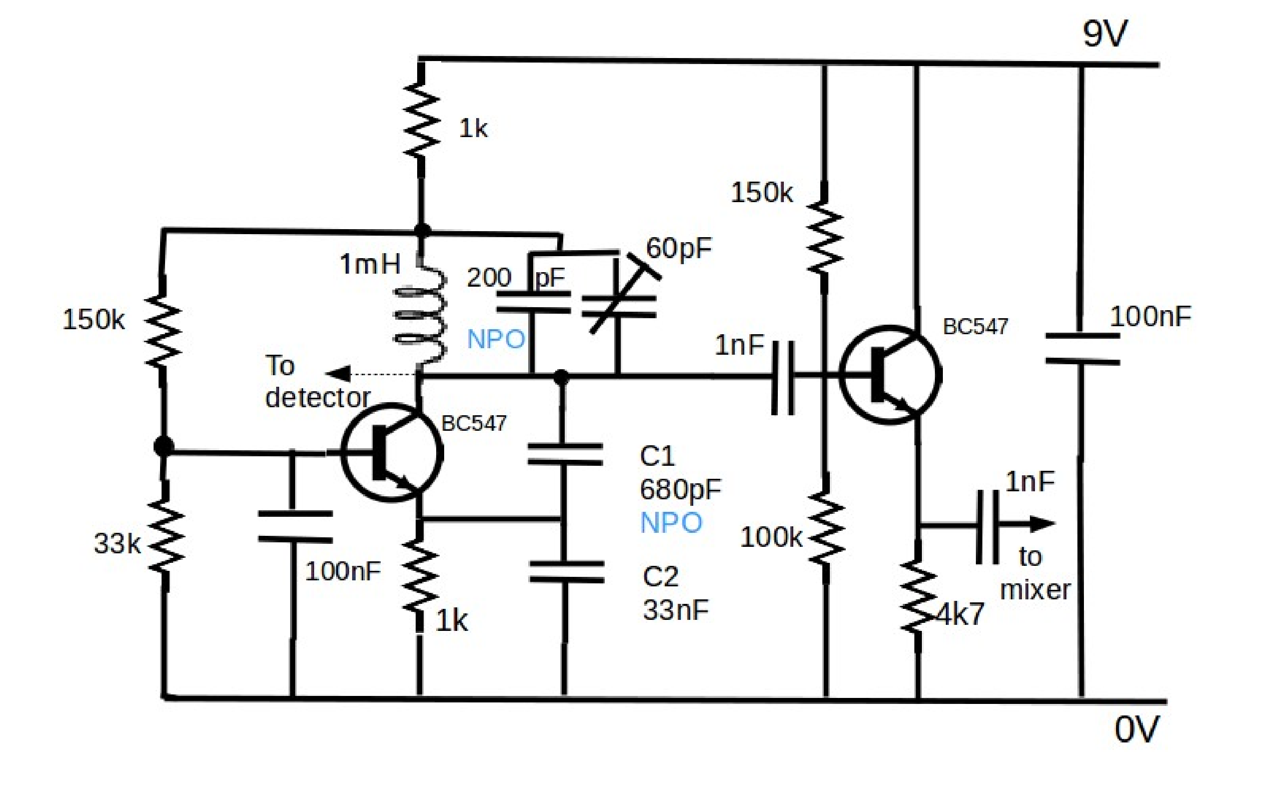

Fig 4.5.7.1. The pitch variable oscillator.

4.5.8 The pitch proximity detector

The electrical theory of the working of the pitch proxmimity detector is the same as that of the volume proximity detector: it is a series resonant circuit very sensitive to added capacitance. The requirements as perceived by a musician are, however, rather different. Musicians like instruments that have a linear relationship between hand position and what musicians call 'pitch'. A piano keyboard is an excellent example of this. The distance along the keyboard between each 'A' note and the 'A' note of the octave above is exactlty the same, and the distance by which one has to move a finger to change from 'A' to 'B' is the same in the top octave as in the bottom one.

Pitch is related to frequency, but the relationship is logarithmic rather than proportional (Fig 4.5.8.1).

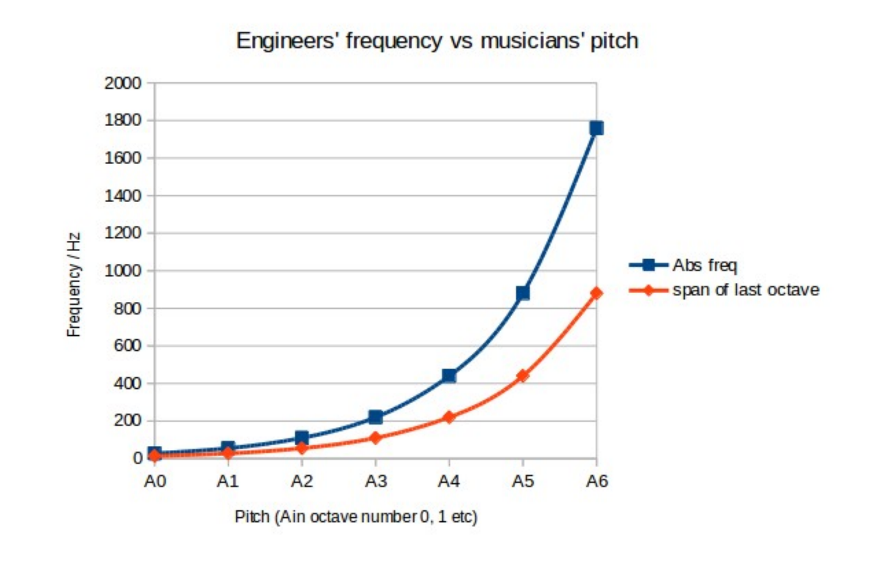

Fig 4.5.8.1: the non-linear relationship between what musicians call pitch and real frequency. The blue line shows the real frequencies of the pitch of 'A' in different octaves. Musicians perveive the 'steps' between the 'A's as the same (' 1 octave'), yet the step in actual frequency (red line) is different every time.

Musicians will want the relationship between hand distance from the pitch half-capacitor and pitch to look like this (Fig 4.5.8.2);

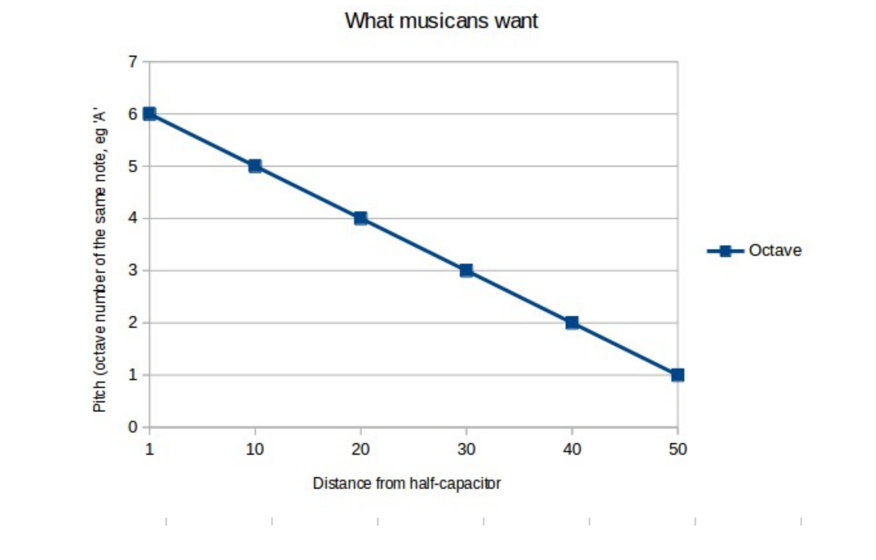

Fig 4.5.8.2: the relationship between musical pitch and distance desired by a musician.

It follows, by combining the two graphs, that the relationship between distance and frequency that the builder has to provide should be as close as possible to this (Fig 4.5.8.3):

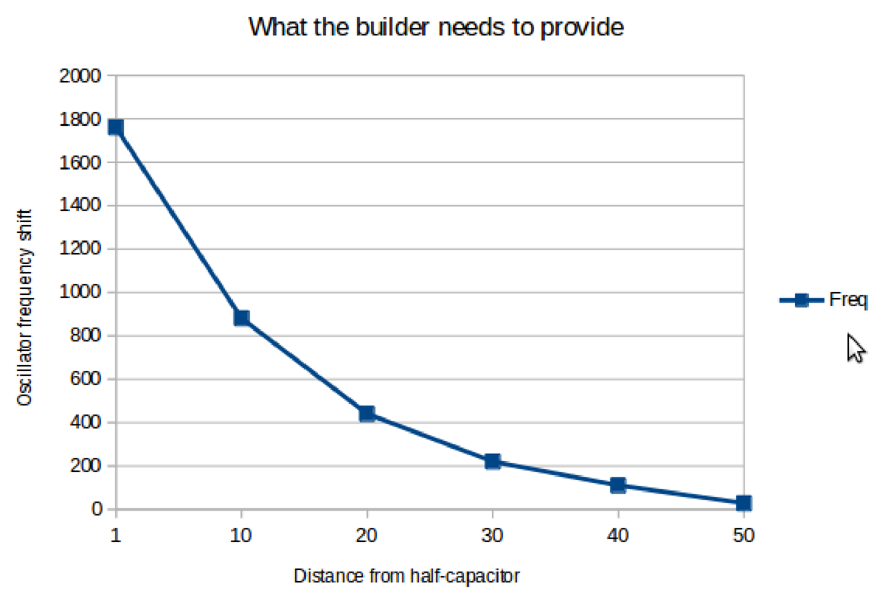

Fig 4.5.8.3: The shape of the frequency-distance relationship needed to produce the distance-pitch relationship desired by the musician. (The line should be a smooth curve, not a series of straight connections between dots, of course: this was plotted at low resolution).

This is where the capacitance-distance relationships of Fig 4.5.5.1 come in useful. Resonant frequency is a function of the square root of capacitance ($f = 1/ 2\pi\sqrt{(LC)}$), so the change with resonant frequency for loop and rod antennae look like this (Fig 4.5.5.2):

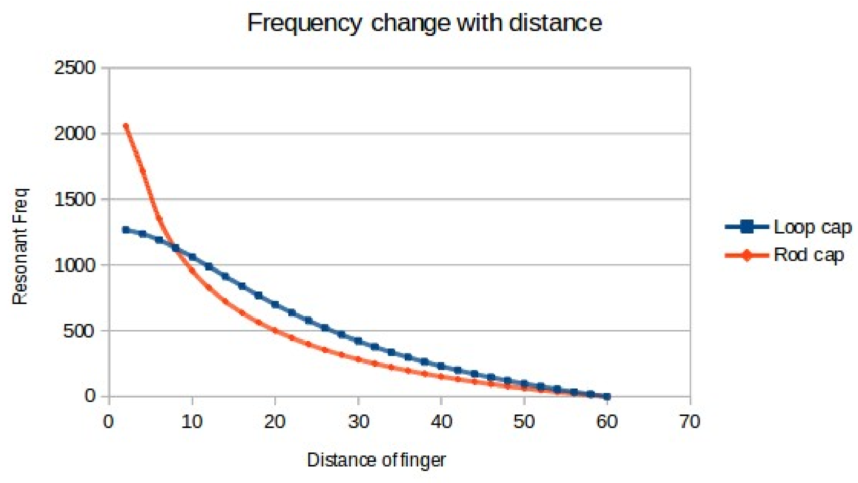

Fig 4.5.5.2: Frequency change with distance for a loop (blue) or rod (red) half-capacitor.

Clearly, the rod is closer to the shape that we want. In fact, if the predictions for the rod (red) are plotted on the same graph as the desired response (blue), we see this (Fig 4.5.5.3):

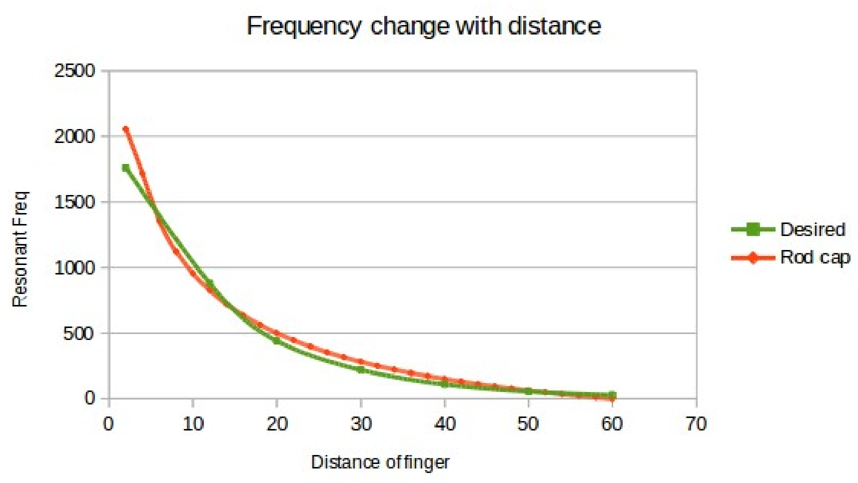

Fig 4.5.5.3: The relationship between distance and frequency shift that is desired (green line) and predicted by using a rod half-capacitor (red). These calculations assume that the correct inductance has been added (realistic) and there are no extraneous capacitances (less so).

This, of course, is what Терме́н must have worked out all those decades ago, without the benefit of computer modelling.

The capacitance of the no-hands rod is similar to that of the loop (of the order of 10pF). For resonance at 190kHz, total inductance can therefore be expected to be about 70mH, although in practice it will probably be less due to stray capacitance that will already pull the resonant frequency lower than calculated on the asusmption that the only capacitance from the rod.

4.5.9 The multiplicative mixer

At the heart of the Termenvox is a multiplicative mixer which, when fed with the signals from the pitch oscillator and pitch reference oscillator, produces an output at a frequency that is the differnce between those input frequencies. The phrase 'multiplicative' is being stressed because the word 'mixer' along means different things to audio engineers and radio engineers and, as the Termenvox uses both radio and audio, just using the word 'mixer' would be ambiguous.

To an audio engineer, 'mixer' usually means an *additive* device: an additive mixer simply adds the instantaneous currents of two input signals to create a combined signal at the output. This type of circuit is used, for example, in a studio mixing desk or a karaoke machine to combine a vocals track with a backing track. Designers of additive mixers take great care that simple addition is all that takes place.

To a radio engineer, 'mixer' usually means a *multiplicative* device. High school trigonometric algebra includes an interesting result of multiplying the sines of two different angles together. In general terms,

$$\sin(A) \times \sin(B) = 0.5\ \{ \cos(A-B) - \cos(A+B)\}$$

If we replace the simple angle terms A and B with terms more relevant to description of a wave, in which instantaneous amplitude is sin ($2\pi ft$), then we can rewrite the above in terms of two input frequencies f_1 and f_2;

$$\sin(2\pi f_1 t) \times \sin(2\pi f_2 t) = 0.5 \cos([2\pi f_1 - 2\pi f_2]t) - \cos([2\pi f_1 + 2\pi f_2]t)$$

In other words, **if we can arrange for a circuit to multiply the input waves of f_1 and f_2 , it will generate two new waves of frequencies $|(f_1-f_2)|$ and (f_1+f_2)**.

Multiplication can be achieved by amplitude modulation. Although AM is not normally thought of as multiplication, it *is*, of course: in a classical AM broadcast transmitter, the instant amplitude of the carrier wave is a product of the unmodulated value of the carrier wave and the value of the modulating wave, eg the radio announcer's voice. That is why the modulating wave (the voice) causes the production of side-bands on the side of the carrier: these are the sum and priduct frequencies.

Many amplitude modulator circuits have been developed over the years. A simple one imposes the carrier wave on the emitter of a transistor, and the modulating wave on the base, overall current in the collector-base path being a product of the two. We can use this for the Termenvox. Two signals will appear at the transistor's collector, the sum and the difference frequencies. The sum frequency, which in this case will be approx 190kHz + (190-minus-a-bit kHz) ≈ 380kHz, is up on the MF radio frequency range, and easily removed by running it to ground through the capacitor of a simple R-C low-pass filter network, leaving just the audio-range difference frequency.

The circuit is shown in fig 4.5.9.1. The potential divider on the input (upstream of the 1nF blocking capacitor feeding the transistor's base) is present just to reduce the amplitude coming from the pitch reference oscillator (only a small signal is needed at the base of the transistor). The values of the blocking capacitors are not critical, and were chosen to give obnly modest capacitive reactance (1nF gives about -840$j\Omega$ at 190kHz); they do not have to be particularly stable. The 1k, 4.7nF R-C circuit that follows the collector forms the low-pass filter (cutoff frequency = $1/2\pi RC \approx 34$kHz, easily high enough to allow the highest musical pitches but 10x smaller than the sum frequency we are wanting to reject). Again, for basic operation of the Termenvox, these values are not critical – just place the cut-off somewhere above 20kHz and well below 190kHz.

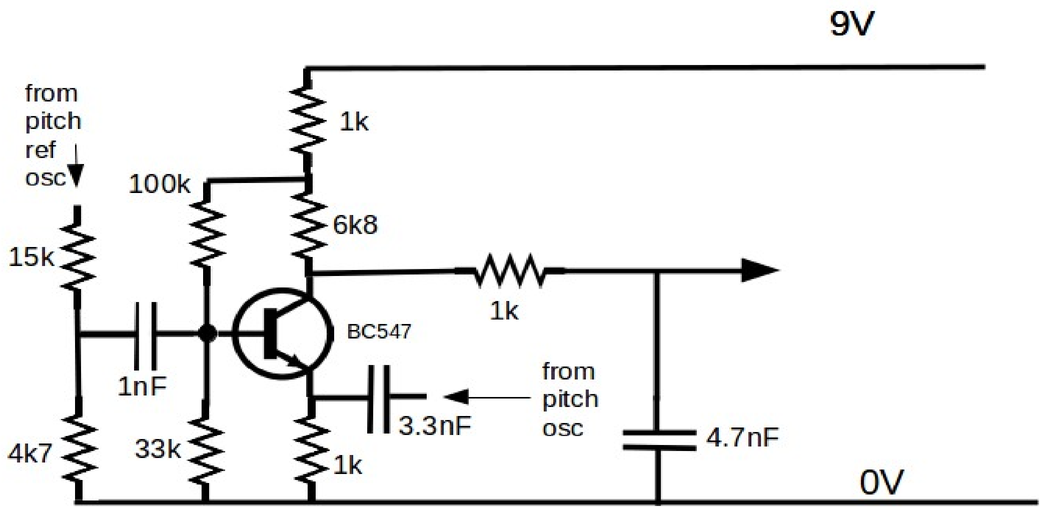

Fig 4.5.9.1: the non-linear mixer. The output signal goes to the voltage-controlled attenuator module.

There may be one reason to have a larger value for RC and to place the cut-off in the low thousands of Hz, depending on taste. Some instruments, including the RCA Theramin, seem to have a pitch-dependent timbre, low frequencies sounding richer and rougher than the pure, sinusoidal high frequencies. This feect is sometimes ascribed to interaction between the pitch oscillator and pitch refernece oscillator (and Терме́н did build a variable weak coupling between them into the special instrument he built for Clara Rockmore, though she usually played with this coupling control at its weakest setting). Another, simpler explanation might be that there was a rather low cut-off frequency downstream of the mixer, so that while harmonics of low notes would get through, harmonics of high notes (which would themselves be much higher) would not, leaving only a fairly pure sine wave. The only copies of the RCA circuit diagram I can find do not have all component values marked, so it is difficult to judge where the cut-off frequency would have been.

4.6 Testing and coarse tuning.

This section assumes that each of the modules has already been tested in isolation.

It is much easier to deal with the pitch and the volume sections separately so begin by switching the tune switch to the 'tune' position. Plug the pitch half-capacitor in and, with no hand near it, and nothing else near it either, use the trimmer on the pitch oscillator to set the frequency to about 190kHz (measure the frequency downstream of the buffer amplifier, so that the capacitance of the frequency meter leads and input stage do not 'pull' the oscillator). Move a hand near the pitch half-capacitor, and observe how much and how linearly the frequency drops. If the frequency rises instead of dropping, a lot more inductance has to be added to the pitch proximity detector circuit. If the frequency drops significantly only when the hand is very close, try adding a modest extra inductance. If the frequency falls when the hand is distant and the oscillation dies altogether when the hand is close, reduce the inductance. This phase has to be empirical, as even small amounts of extraneous capacitive reactance, dependent on the exact box, plugs, wiring routes used, will make a large difference to the resonant frequency of the proximity detector. Unfortunately, having made a change to the inductance, you may need to re-tune the oscillator itself back to 190kHz. This phase may take time, but has to be done only once.

Next, set the coarse pitch null control in the middle of its range, and use the trimmer to tune the pitch reference oscillator to the same frequency of the 'no hands' pitch oscillator. While a frequency counter can help, the final tuning is most easily done by ear (when the frequencies are close, the audio difference frequency will be heard). If nothing is heard when the frequencies of the oscillators are between 200 and 10,000 Hz of each other, something is wrong. Given that the audio amplifier will have been tested already, the most likely culprit is the mixer.

With sound coming through, finish the tuning for a null (with no hands near the pitch rod), and then move the coarse pitch null control so that a sound appears again. Then switch the tuning control to 'play'. Plug the volume loop in, keeping hands, tables etc well away from it, set the volume null control to its mid position, and use the trimmer to set the volume oscillator to 454kHz. You should now hear the note. Move a hand to the loop; the frequency should drop to less than 453 kHz. If it rises, add a lot more inductance to the volume proximity detector. If it drops but too little or too late, add modest amounts of extra inductance, again empirically. With the inductance OK, stand in a playing position and move a hand towards the volume loop. The volume should fade steadily until a rapid drop at the end. If it fades very suddenly from loud to silent, especially if there is a 'clunk' or 'whump', the most likely problem is too high a Q factor in the frequency-to-voltage converter. Increase the resistance on the ground line of the resonator or, if you used the idea of a rheostat on the front panel, adjust it to the higher resistance (less sensitivity) end.

If the tone of the instrument is too rough, (has too many harmonics), check that the output signals entering the multiplicative mixer are not clipped (they have no flat tops or bottoms when viewed on an oscilloscope). If any do, find out if the problem is in the oscillator itself or the buffer amplifier, and adjust gain/ coupling accordingly (though there should be no problem – there was none in my own build). If the signals entering the mixer are pure, try reducing their amplitude as they enter the mixer to reduce any tendency of the mixer to clip (the signal entering the base is the most likely one to require this treatment). If there is no serious clipping, a simple approach to reducing harmonics may be to increase the value of RC at the end of the mixer, to reduce the cutoff frequency. This has the interesting effect of making low notes richer and high notes purer.

5. Tuning steps for the musician

The instrument can be set up with the follwing steps (it is not that complicated – there are many steps below just to make sure that nothing is ambiguous);

- Place the instrument securely on a stand or table before attaching the vulnerable pitch rod and colume loop.
- Turn the volume control down (not essential, but is prevents startling noises on power-up).
- Turn the fine pitch-null control and the fine volume null controls to their middle positions.
- Plug the power supply lead in to the Termenvox; ensure it is not a trip hazard.
- Plug the pitch rod into the top of the instrument, and the volume loop into the side.
- Plug the power supply into a mains socket and switch on.
- Wait 30 seconds for everything to stabilize.

- Switch the tuning switch to the 'tune' position.
- Gently turn the volume control up about half way. There may or may not be a sound.
- With hands, head etc. well away from the pitch rod, turn the coarse pitch-null control until you hear a mid-pitch sound (if you can already hear something, skip this step).
- Switch the tuning switch to the 'play position'.
- If you cannot hear anything, alter the volume coarse tune until you can.
- Move your hand close to the loop. The volume should diminish to nothing. If it does not, alter the coarse tuning until it does. You are aiming for sound when your hand is far from the loop, and silence when it is close. Use the fine control for final adjustments. The volume control is now set up.
- With hands away from the volume loop, turn the coarse pitch-null control in the right direction for the pitch you hear to fall (this could be either way). Keep turning until the pitch is either inaudibly low or at least very low.
- If you kept overshooting the null with the coarse pitch-null control, use the fine pitch-null control to set the pitch inaudibly low.
- Move your hand towards the pitch rod: You should hear a low note rising higher as you move your hand nearer to the rod (try not to touch the rod: it won't hurt you, but you may 'stall' the instrument).
- The instrument is now ready to play.

Common problems:

* After the pitch tuning is complete, the volume control will not silence the note wherever the volume null control is set and whether or not the hand is near the loop.

→ The tune switch is probably still set to 'tune' not 'play'

* The pitch control works 'backwards', pitch being high with no hand and falling when the hand approaches.

→ The coarse pitch-null control is in the wrong place by a large margin.

* The volume control works 'backwards', being quiet with no hands and loud with a hand near the loop.

→ The volume-null control is in the wrong place by a large margin.

6. Playing the Termenvox

The Termenvox is usually played from a standing position: ensure that hair and clothing is not so loose that it will move 'randomly' near the loop or rod. The descriptions below, which provide only elementary information, are for a right-handed instrument: reverse them for a left-handed one.

Volume: with the left hand well away from the volume loop, move your right hand towards trhe pitch rod. You should hear a rising note – stop when the note is comfortable to listen to (too close might be annoyingly high). Now, leaving your right hand where it is, slowly bring your left hand, held flat, down towards the loop. The volume should diminish, and cut out completely when your hand is close to the loop. That is one way of manipulating volume (and is slow; useful for legato). Try keeping your left hand at the same height, and slide it away from the loop, horizontally: the volume will rise again. If you 'wag' your hand to and fro in that horizontal plane, you can make relatively rapid changes. That is one other way of modulating volume. Now move your hand awa again, and point your fingers (all, or index and second) downwards. Move them down into the loop; the volume will die away. Lifing your hand from the wrist so that the fingers 'jump' up and down, in and out of the loop, can provide very rapid changes in volume, ideal for staccato playing. Once you have the feel of this, move your right hand slowly away from the rod as you 'bounce' your left hand's fingers into and out of the loop. What you will hear will be nothing like a proper scale, but it will give an idea of how moving the left hand can break a 'swoop; of pitch into discrete 'note-like' pitches (which will be real notes when you have the hang of the pitch system).

Pitch: with the left hand away from the volume loop, hold your right hand, with thumb uppermost and fingers relaxed (relaxed fingers will curl so that the finger tips point at about 90 degrees to the palm). Move your hand towards the rod so that it is on the outside of your fingers, as if you were going to punch it. There will be a tone. Now experiment with opening moving one, then two, then three fingers, so that they approach the rod. This will give you fine control of pitch. Course control comes from moving your whole hand. You can do this from the wrist (fast, fine), the elbow of the shoulder. Experiment.

Vibrato: With your right hand held vertically with fingers straight out, bring your right hand close enough to the rod to hear a mid-range note, then use your writst to pivot your hand very rapidly back and forward a few degrees, about the axis of your index finger (that is the most natural pivoting movement). The lower part of your hand will change its distance from the rod, and a vibrato tone will be heard.

7. Things I might do differently if I build another one.

While building this Termenvox, a few ideas occurred to me later than I could easily incorporate them. Some of these ideas were for a completely different way of making a hands-free instrument (and these will probably appear on this website in a few months). Others were for small changes of a still classical Termenvox:

- I would not bother to build the fine pitch control: the 'coarse' control is easily fine enough for the musician to tune the pitch part fo the instrument accurately and the fine control remains untouched.
- I might include a swtich that hands the control of the VCA module either to the existing volume module, or to an external device such as a foot pedal so that the musician can play this instrument either as a classical Termenvox, using the internal volume module, or like the earlier Etherphone instrument.
- Having built it, the oscillator-then-resonant filter arrangement for volume control feels inelegant (it works, but uses more components than should be needed for a simple job). A grid dip oscillator (strictly 'gate dip', using a FET, but the old thermatron* name has stuck) would probably be more elegant, and more similar to what Терме́н did himself. Some work would be needed, though, to ensure that the sensitivity was correct. For anyone who just wants a working instrument without doing their own research and development work, it would be easier to follow what I did here.
- I might follow my own suggestion, made in Fig 4.5.3.3.5, and make the Q-spoiling resistor in the ground line of the resonator in the VCA module adjustable. On reflection, routing the line to and from a rheostat (5k?) on the front panel would be asking for trouble from stray inductances and capacitances. I would therefore probably use a resistive optical isolator, the light-dependent resistor part being in the ground connetcion of the resonator, and a variable dc current driving the LED part of the optical isolator via a rheostat or via a transistor, the base of which would be attached to a potentiometer on the front panel. The transistor idea would be better: the potentiometer could be connected to fixed resistors so that the whole travel of the potentiometer altered the LED brightness in the optical isolator to give a range of resistance in the light-dependent resistor from about 10Ω to about 10kΩ; a logarithmic potentiometer would probably be best. To reach down to the 10Ω range, it may be necessary to use several optical isolators in parallel (their output resistors may not go down to 10Ω individually, even at full brightness).
- It may have been useful to build a fixed frequnecy audio signal generator into the instrument, to help setting up the volume control system. That way, pitch can be set up even if the volume section is not known to be set up properly (the tune switch does this), and the volume can be set up even if the pitch is not known to be set up properly.

* 'Thermatron' is a recent term for what were called 'valves' in the UK and 'tubes' in the USA. It was made popular by the superb book *Hollow State Design*, by Grayson Evans (TA2ZGE), 2013, ISBN 978-1-300-96521-3

8. References

***DTDFCM**: Marston RM (1991) Diode, transistor and FET circuits manual. Newnes/ Butterworht-Heinermann (London)/ BPB Publications (New Delhi). ISBN 81-7029-410-X

E&M: Duffin WJ (1980) Electricity and Magnetism, 3rd Edition, McGraw-Hill,
ISBN 0-07-084111-X

EMRFD: Hayward W, W7ZOI; Campbell R, KK7B; Larkin B W7PUA (2012) Experimental Methods in RF Design (Revised first edition). ARRL Press.

QRPB: Dobbs G, G3RJV (2012) QRP basics, 2nd Edition. Radio Society of Great Britain. ISBN 9781-9050-8684-9