Constructing A Driver Transformer For Thomas Vox Amplifiers

Table of Contents

Transformer design and rewinding are very detailed skills. The information presented here is necessarily a very sketchy and pared-down view of these crafts. The information has been cut back to just what is needed for this specific task.

Much of this information is taken from my rewriting and editing posts I've made over time on Vox-related forums. I hope I've corrected the mistakes and made it more consistent, as well as updating it with information I've picked up since those posts.

Why are we doing this?

The Thomas Organ Vox amplifiers were never common, and they're becoming downright rare. When these things break, they can be hard to find parts for.

About the hardest things to find a suitable substitute for are the driver transformers. These amps were designed at a time when the more modern complementary or quasi-complementary output stages for Class AB amplifiers were not yet in common use. Instead, they used a two-transistor amplifier to drive the primary of a driver transformer, which then ran the output transistors. This simplified things a bit, and was fairly well known and tested at the time. But time marches on, and the transformer driver became quickly obsolete as cheaper, easier solid state devices and techniques became available.

Luckily, the driver transformer is generally more durable than the output devices in these amplifiers. But sometimes an output failure can fry one. When that happens, you have only a few options, one being to junk the amplifier. I hate that to happen, so I started digging for other options.

Something that's not an option is buying another driver transformer. There aren't any that aren't inside amplifiers. Cannibalization works, but still leaves one dead amp. Another is to send the dead transformer to a transformer re-winder to un-wind, and reverse engineer, so they can wind you a new works-alike. This is workable, but very expensive.

I've worked on a third option. That's to gather enough electronic intelligence from testing on a working driver transformer to create a recipe for making a functionally equivalent replacement out of another transformer. The idea is to find another transformer of about the same size and shape as the originals were, take it apart, and wind new windings on it that make it work the same as the original Vox driver transformers.

This description is wordy. If you know something about transformers already and just want to get on with it, you can skip to Appendix A for a much more concise plan of attack.

The Originals

The original Thomas Vox driver transformers operated like the output transformers of small audio power amplifiers themselves. It took me quite some time to come to this realization, as there is no good technical description of how the driver circuits work that I've ever been able to find. And that includes some years of poring over musty textbooks and circuit design manuals.

They must deliver a notable amount of power to the output stage to work properly, and in addition deliver it to two separate outputs.

The Circuits

The circuit the driver transformer works in was referred to – when it's referred to at all – as a Class AB stacked output or sometimes “totem pole” output, the idea being that two almost identical stages are stacked on on top of one another. Each section of the output delivers one polarity of the output signal. The sections of the totem pole are driven in the proper phase by the driver transformer. The transformer's job is to make the signal polarities work out so the two outputs work in tandem, and to allow the DC stacking of the top and bottom sections to let the outputs reassemble an output.

The quirk is that the two secondaries drive the positive-side (top, or upper) output devices and negative side (bottom, or lower) output devices alternately, the junctions of the output devices themselves picking which one is active by which polarity of voltage on its secondary makes it conduct. The other side is turned off by the opposite polarity on its secondary. So the two secondaries taken together put out the entire signal from the transformer into a load that's similar to an audio speaker.

After pondering this circuit at times over literally years, it became clear to me why it was forgotten so quickly. It's expensive, as witness this transformer we're worried about here; it's not very efficient, and wastes a lot of power compared to the more modern complementary and quasi-complementary output stages we use today. And it's hard to protect against load faults.

It was only this realization that finally got me to where I could analyze the circuit and begin to understand it. I'll add an appendix with some details of the circuit's design if I get time.

What's good about it from a musician's standpoint is that it sounds better than high-feedback solid state designs. It has some internal distortion, but what distortion it has is less objectionable than many solid state designs.

The Original Transformers

The biggest of the driver transformers was the one in the Beatle amplifiers, which put out 120Wrms into 2 ohms. These driver transformers put out several watts into a load that's very similar to an 8 ohm speaker.

As I mentioned, the totem pole circuit is inefficient. The driver transformer puts out nearly 8W of audio power in the Beatle amps to drive the four output transistors. It's intended to do this down to bass-guitar frequencies of 40Hz. The driver circuit runs the driver transformer in single ended mode, so it's using only a fraction of the energy storage abilities of the transformer core.

Those taken together make the transformer big for the power it handles. The core and windings are as large as a 60-Hz power transformer rated for about 20-24W of AC output power.

But the good news is that these relics used a lamination size and style that is still widely available; and copper wire has not gone out of style. We can find modern substitutes for the parts of the transformer and wind something that will act much the same.

The driver transformers scale down as the output power of the amplifier they work in scales down. The Royal Guardsman driver is a bit smaller than the Beatle, the Buckingham, smaller; the Berkeleys and Cambridge smaller still.

Since we have the luxury of not having to meet a manufacturing budget, and since the Beatle-size transformer will work in all the circuits, we only have to get one design working. I believe this size will drive any of the circuits well enough. This is good – smaller is more difficult to make in transformers.

Overview of the process

For making your own clone driver, you're going to have to do the following:

- gather materials, especially a donor transformer; don't start this before you have your donor transformer in hand
- disassemble the donor transformer
- gather tools

- set up to wind coils
- wind the coils
- stack the laminations and reassemble the transformer
- test the result
- rewind if necessary
- pot the resulting device

Obtaining materials

Donor transformer

The original Beatle driver was made with 75EI scrap-less laminations (that means a standard shape with the middle of the E being 0.75” wide) stacked 0.75” high.

Something very near that size would be really handy. And after a little searching, I found several. The Triad/Magnetek F-45X is a 1” stack of 75EI. It's in the picture at lower left. There is a similar Hammond transformer, shown at lower right.

For the purposes of finding a donor, power line transformers are sized in a way that the overall core and windings are proportional to their output power. A 24V/1A transformer is about perfect. So would a 12V / 2A transformer be, or an 18V / 1.5A transformer. 24 volts times amps seems to be right. The pictures I could find showed the Triad with non-split bobbin, but the part I actually got had a bobbin like the Hammond.

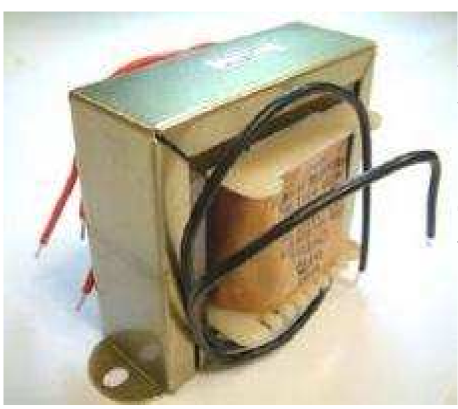

Transformer materials

The donor transformer supplies the iron cores, the bobbin, and the frame that holds it all together. The wire inside the donor transformer will probably be useless for the final product, because (1) the wire sizes are critical to making all this fit and (2) you're likely to scrape the insulation or kink the wire in disassembling the donor. So trash the old wire, or if you're thrifty, save it for another project.

In addition to the donor transformer, you're going to need

- magnet wire
- layer insulation, such as mylar film.
- high temperature tape

– wires for leads outside the transformer
– solder
– nylon fishing line to use as a spacer on the ends of the winding; 15 pound rated is about right.
– some way to protect the outside of the windings
– potting material; short-oil varnish; phenolic vs alkyd; not polyurethane[1]
– high temperature tape to hold things together
– stranded wire for the external leads
– optionally, fiberglass tape and/or transformer varnish for a bang-up professional job
– some way to hold the bobbin while you wind the new windings on the bobbin
– patience; lots of patience.

Tools

You'll need at least the following tools

– X-Acto knife
– small flat screwdriver
– diagonal cutters
– file
– scissors
– small hammer or wooden mallet
– soldering iron
– Sandpaper for stripping magnet wire
– possibly woodworking tools to make a bobbin form, and/or coil winding setup

Winding setup

While it is theoretically possible to wind one of these things up with nothing more than your hands and some tape, I would not want to try it. That job needs about twelve hands. To do a reasonable job, you're going to need a way to turn the bobbin on its axis while you feed wire to it in a neat and orderly fashion.

The simplest way I've seen is to cut a block of wood that just barely fits through the hole in the bobbin, and then to fix that block so it can rotate on an axis. You also need a brake or stop to keep the windings from all spilling off if you relax your hold on the bobbin. I'll include some sketches as suggestions in an appendix.

You'll also need some way to count turns. Counting in your head is a poor choice, as there are a lot of turns, and it's easy to lose count. The number of turns on this thing is very important, so you need to figure some way to do it right.

I found a Dollar Tree store and on a whim went in. They had both four-function and scientific calculators for $1.00 each. All modern calculators do chained operations, so that if you enter, say, 35,

1 Varnish is the chemical result of cooking an oil and a resin till they link up, suspended in a solvent. The most common home-supply-center varnish is polyurethane-based, but this is not good for transformer use. You want alkyd or phenolic resin-based varnish, and a “short-oil” kind -cooked with a minimal amount of oil.

and then + 1 =, you get 36; if you simply hit equal again, it shows 37, repeating the + 1 to itself internally. So you can set this up to count by 1 if you can arrange a way for the = key to be hit each turn. That can be mechanical or electronic. I'll give an example in an appendix.

Even if you have to count off ten turns, then make a mark on a piece of paper, figure out some way to keep track of turns other than a count in your head.

Arrange the wire so it spools off easily, but with a slight drag, again to keep it from spilling off if not held.

Potting setup

At the end of the winding, when it's all taped up and done, you'll have to choose one of four ways to pot it:

1. don't pot it at all
2. pot it in wax
3. pot it in varnish
4. Figure out something different on your own.

Transformers operate more quietly and cooler if you fill all the air spaces inside it with some kind of material to act as a mild glue, much like the laminations had when you first took them out. It is possible to use the transformer without this, but the transformer will live longer and run cooler if it's potted.

There are are two ways that this can be done that I can recommend from my reading. One is to pot it in beeswax or a mixture of beeswax and paraffin wax, the other is to dunk it in varnish and then heat-dry it. It's hard to say which is better.

Salvaging A Donor Transformer

Don't start the project until you find a suitable donor!

The Triad F-45X would also make a good donor, and Mouser has stock on these for $11.00 roughly at the time I'm writing this. The pictures are of the F-45X I bought for my prototype.

Even cheaper, get out your ruler and/or dial calipers and head to a surplus store. You want a transformer with a body of iron laminations which measures 0.75 or slightly more across the thickness of the central stack, 1.875" at the raw iron laminations the short way, about 1.9" to 2.0" overall if it's wrapped in a sheet metal frame as many of these transformers are, and about 2.3" to 2.5" the long way, again if wrapped in a sheet-metal mounting frame.

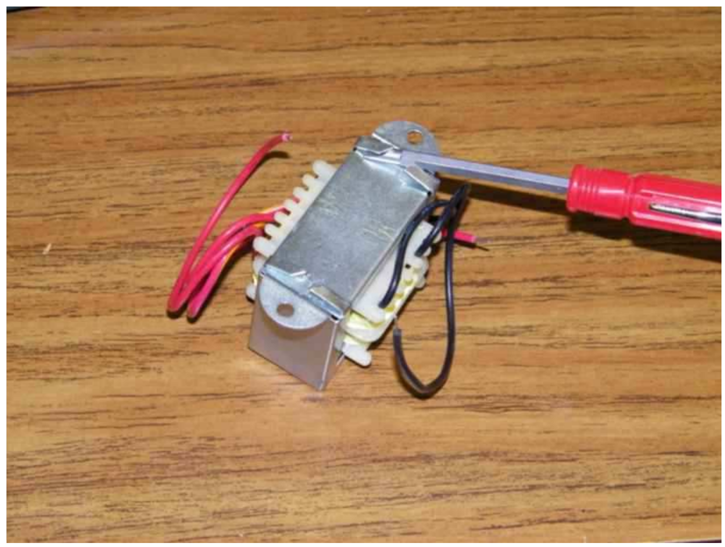

If you're looking for surplus donors, look for and pass on any donor which has a bead of welding across the iron core laminations. This can't be used because it cannot be disassembled.

The top picture is my donor, the Triad F-45X, before work begins on it.

The middle picture shows how to begin dis-assembly Gently pry up the folded metal tabs holding the bottom plate on. Your donor may or may not have a bottom plate.

The bottom picture shows the tabs fully bent up. Do as little bending as possible on the tabs, because you will have to bend them back down to complete the reassembly of the transformer, and you don't want them breaking off when you do.

On the top at right, the bottom plate has been removed.

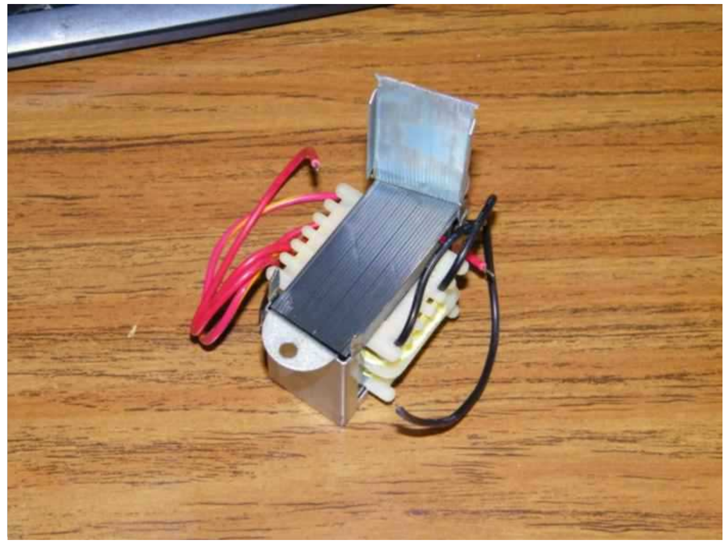

The middle picture shows the frame loosened and being removed from the laminations.

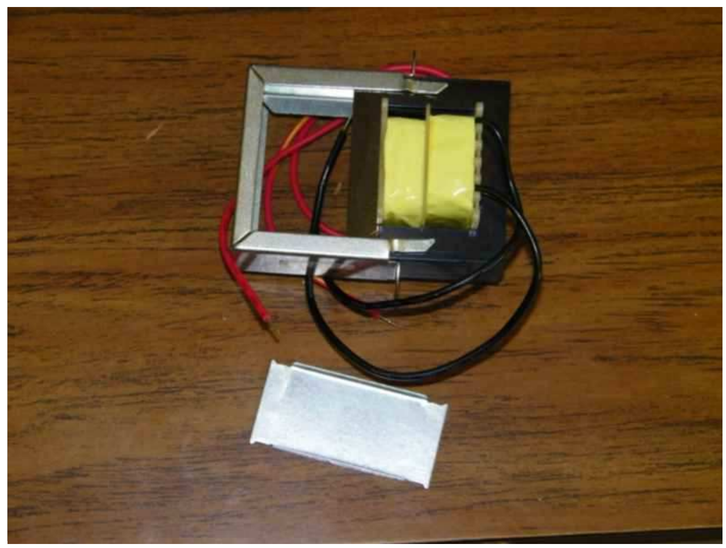

You'll notice that the parts are stuck together with varnish that holds them in place until it's broken loose. You may need to tap pieces to break the varnish loose all through the dis-assembly process. Do this gently, as we need the parts under that varnish.

The bottom picture shows the first “I” lamination removed. I used a small pocketknife as shown for this one, but I quickly found that the pocketknife blade was too thick and could easily bend the laminations when I tried it on the first “E” lamination. So I shifted over to using the X-Acto knife for all further laminations.

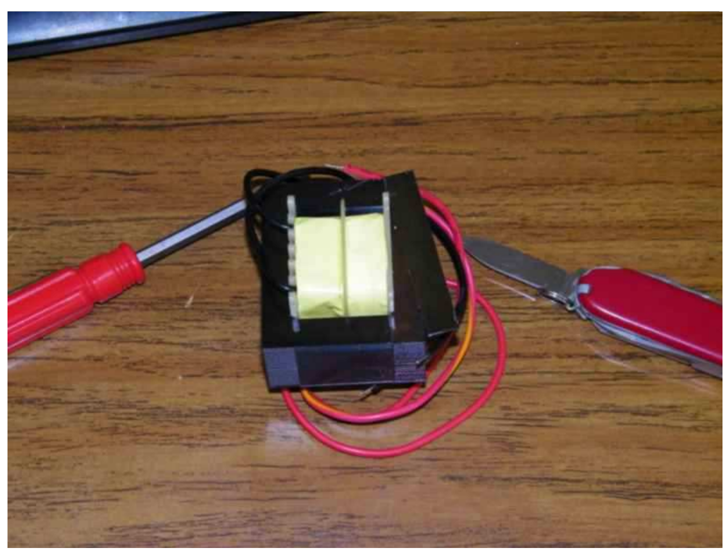

The process is to place the X-Acto blade edge right in the crack between the lamination you're trying to remove and the rest of the stack. Then tap the back of the blade gently to force it into the crack.

This results in an audible “snap” as the varnish gluing the two layers together breaks.

The top picture at right shows where to start with getting the first “E” out. Once the corners of the E lamination have been cracked free, and the edges all around the bobbin have been freed, the work begins.

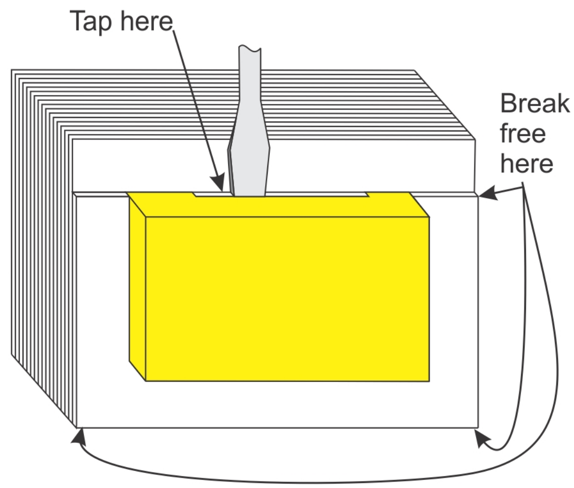

It helps to push the X-Acto blade as far into the middle under the bobbin as possible. You simply have to be patient and work in small steps. You can't afford to break the bobbin. Eventually you will be able to tap on the middle and outer E legs and the lamination will slide out a bit on the bottom.

The picture at middle right is the edge of the first “E” lamination peeking out above the entire stack. It took me about 30 minutes to get from removing the first “I” to this point.

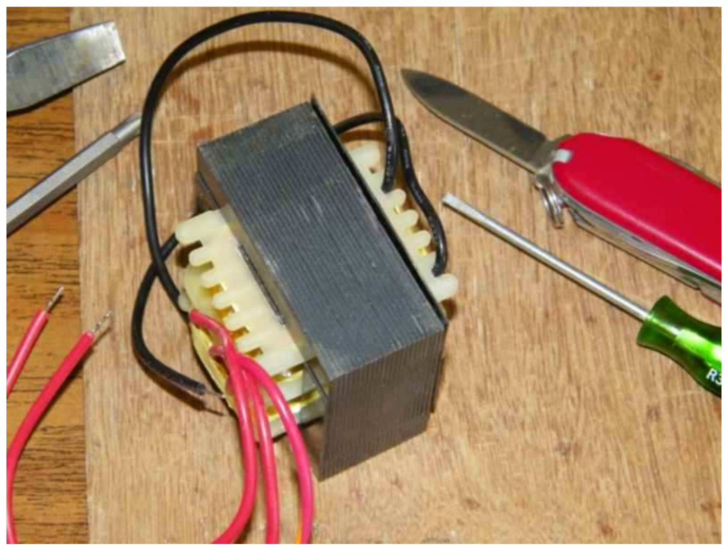

This is because the varnish/glue is also between the laminations on the middle legs of the “E”s and the first one is held in place by being forced into the bobbin as well as the varnish.

Keep working at it and you'll get the first few layers out. When that happens, you're all set up, and eventually they start coming out easily. The picture lower right shows how I started each layer with gentle taps to break the varnish free.

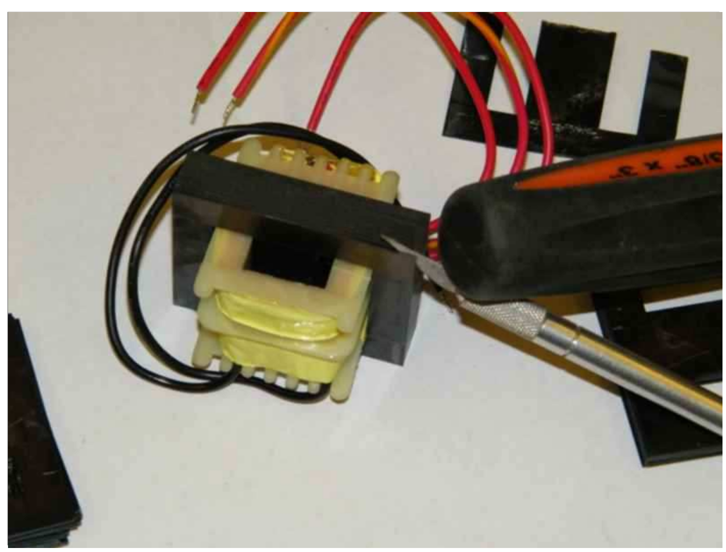

Remember that you're going to have to use these again, so be very careful not to bend them. My donor had been carefully manufactured, because the laminations were quite soft, and if I missed the crack between layers, the X-Acto could actually cut into the iron. They bend easily.

You'll find that the E's are interleaved in alternating directions, then the Is are stuffed into the leftover spaces. Take off an I, then work out an E. Turn the transformer around and do it again til you have them all out.

It is very likely that you will bend up the first E so badly that it will not be usable. That's OK. The donor most likely has more laminations than the original, and this is one place where you just have to get enough iron into the transformer, not just the perfect amount. More is better, and you started with a slight excess. Eventually, you'll get all of the laminations loose and have a nice pile of laminations – and the bent-up first one, as shown at top right.

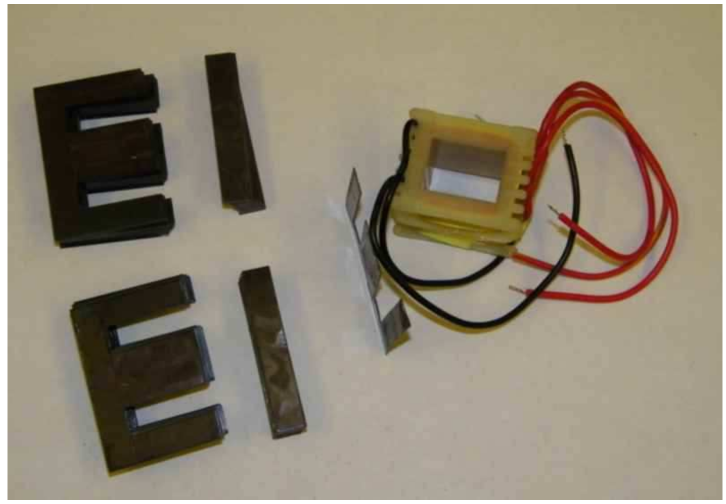

After that, you can start disassembling the bobbin. It is highly likely that your donor will have a nylon bobbin if it was made in the last decade.

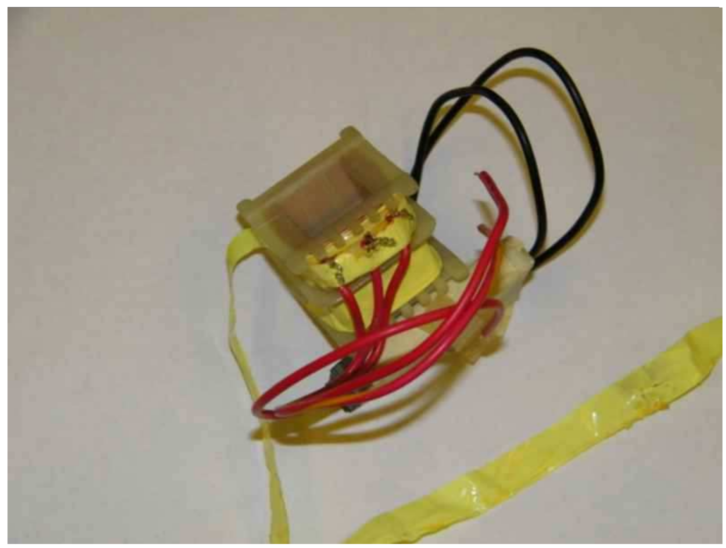

Start by unwrapping the tape which covers one of the windings. In the picture at middle right, I started with the secondaries.

There will be 2-3 layers of outer covering to protect the wires, then the connections of the external wires to the internal magnet wires. In this one, there was a soft plastic pad under the solder joints to keep them from digging into the magnet wire.

The picture at bottom right shows how the magnet wire was attached to the external leads.

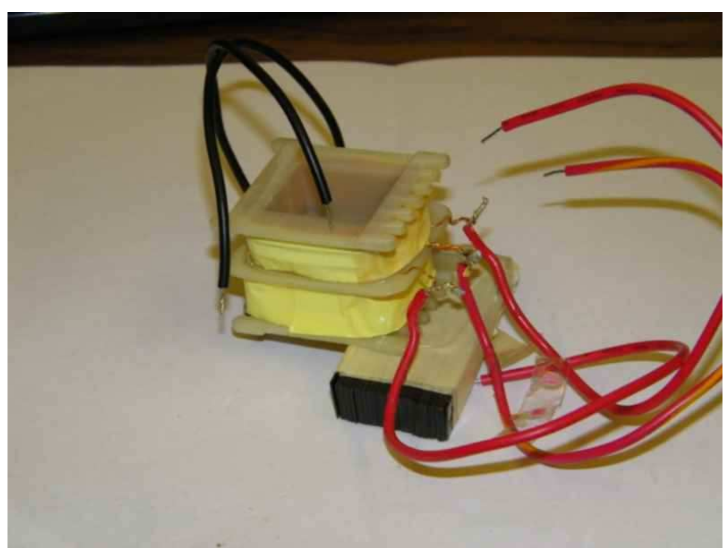

It's important to look at how the windings were constructed, not just rip them out of the bobbin. The way they're put together is a great lesson in how to reassemble your driver transformer. You're going to have to construct something very much like you're taking apart.

Take special note of how the magnet wire is attached to the windings under it. Doing this neatly and well is an important part of getting the wires where they have to go inside the bobbin. How well your driver transformer works depends on how well you can put the wires where you want them in neat, tidy layers.

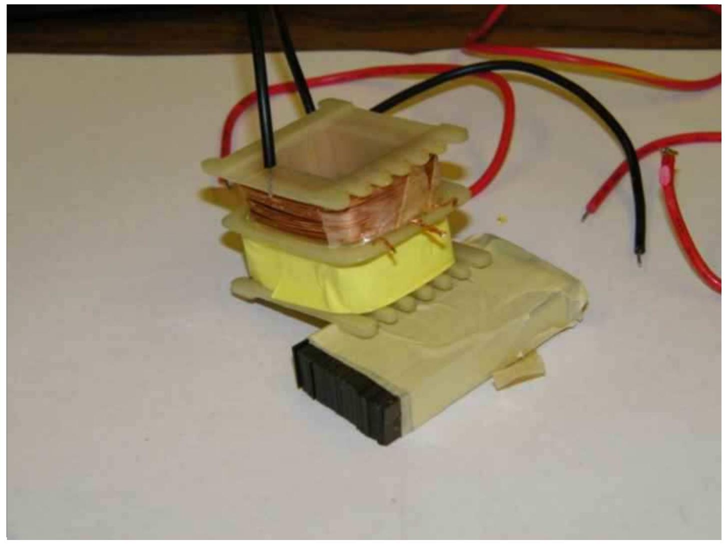

Work patiently and deliberately. If you get frustrated, leave it where it is and go do something else for a while until your mind is clearer and you can focus.

You'll notice that I have wrapped the Es and Is in masking tape. It's likely that I will have to leave this for some days before getting back to it, and I wanted to be sure that they stayed where I could find ail the parts and pick up where I left off.

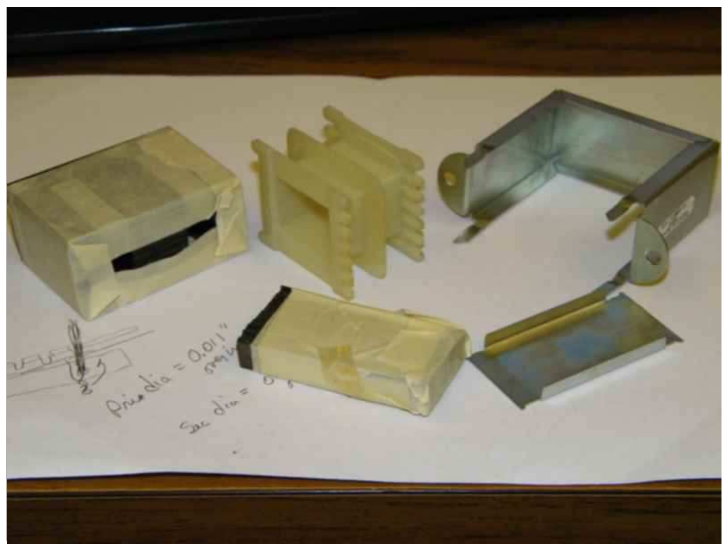

In the picture at middle right, notice that the bobbin has a “wall” in the middle of it. This was used in the original for safety isolation between the hazardous AC mains voltage winding and the secondary windings.

In the case of a driver transformer, we are not using hazardous AC mains voltages on either winding, and the voltages are much lower than the AC mains. We do not need the isolating wall, and it would make the audio performance worse if we left it in.

Cut the wall out of the donor's bobbin if there is one in yours. I nipped small pieces out with diagonal cutters until all that was left was a small ring around the bobbin's middle. Then I used a small file to file away the remaining parts of the wall and make the middle of the bobbin flat and smooth on each surface and all the corners. This will be important when winding a replacement, where the wires have to be in carefully-chosen places relative to one another.

As I mentioned, it's tempting to think that the wire you're removing from the donor could be re-used. It's possible, but this is usually false economy. Even if it's the right size, it is likely to be damaged in removing it from the donor. I tossed all the wire from my donor. The more I thought about there being possibly a shorted turn when I got done with all this work, the more I wanted new wire.

Insulating Materials

The transformer industry set up several insulation classes, based on the max temperature rise. These classes are confusingly labeled Class A, Class B, Class E, Class F, Class H, etc. Not the same as amplifier bias, just the same name. Shortening things up a bit, all the simple, ordinary commercial stuff is designed with Class A insulation, which implies a hottest spot inside the transformer of less than 105C. This insulation class allows the use of things like cardboard tubes, kraft paper layer insulation, and cheap magnet wire. The donor you find will almost certainly be designed with Class A insulation.

I originally thought that any normal household tapes would not be usable. I find I was modestly mistaken. Transformers can get hot inside; for transformer work, you need either polyester or polyimide/Kapton tape rated for some high temperature, over 105C. This can be tricky to find, but you can get it if you google things like polyester tape, polyimide tape, Kapton tape, transformer tape, and so on. It can be expensive. I found reasonably priced transformer tape on ebay. A transformer on a donor bobbin won't need much else.

However, in searching for inter-layer insulation, I stumbled onto the fact that heavy duty packing tape is actually bi-axially oriented polypropylene, and the film won't break down til it's over about 160C. There's no way to tell where the (probably) acrylic adhesive will go, but for inter-layer insulation, packaging tape is probably OK. It won't work for the outside, where the adhesive holding power is important, but for inside layers, it may work.

A deluxe job would use glass cloth tape on the outside, but you may get away with using filament tape here.

The wire insulation is also a component of the insulation system. There are solderable and non-solderable wire insulations. The solderable ones are viewed as “good enough”. I bought the 200C rated non-solderable ones, as I got used to sandpapering off wire insulation back when I was designing transformers. It's your choice on wire-insulation type.

Winding

Sizing the window and wires

Knowing the measured size of the original transformer lets me look up the lamination sizes in a reference book with some accuracy. All normal transformer laminations are so-called scrap-less E-I laminations, and these have a very fixed set of relative dimensions. In particular, the window area to wind wire into is well known.

The simplest approach to sizing wires in a transformer is to note that you'd like no hot spots inside to burn out a small area, which kills it. If you take the area of the winding window and subdivide it on an area basis for half to primary (where all the power comes from) and half to secondaries then equal power density is approximated. You subdivide the secondary window area the same way. This is not perfect, and usually needs tinkering, but it's the way to start.

Notice you don't get to use all the space inside the core window. You have to have a bobbin to insulate the wires from the core, some kind of outer covering to insulate the outside of the wires from the world, and then you can't perfectly fill the remaining area with wire.

I spent a couple of hours, picking a wire size, computing windings per layer, layers, height build up, for primary, then secondaries, and then making adjustments and calculating again.

It came down to six layers of #29AWG for a total of 444 turns, five layers of 30 turns bifilar #27AWG plus one layer of 17 double turns for the secondaries.

If you want to dabble in this a little yourself, I have reproduced some of my calculations in Appendix D, plus an alternate solution. This should enable you to calculate some alternates if you want, or if you have a different core.

Layers and Interleaving

The optimum way to wind this is with three layers of primary, then six layers of secondary, and finally another three layers of primary. This interleaves the secondaries between two halves of the primary, cutting down on leakage inductance.
The secondaries are wound bi-filar, meaning “in two parallel lines”. Literally, you wind on two secondaries side by side at the same time. The secondary wires are so tightly coupled that they have almost no leakage between them, which is a necessity for the way the output stage works.

Winding Tricks

There are a number of tricks to getting wires secured and held in place in a transformer.

Spacers for Neat Layers

The number of turns per layer is not even, and to get nice, neat windings, we must work hard to keep turns from falling off the ends of each layer. Since the ends of the windings have to escape from where they start down in the layers, we have to prevent the layers from getting uneven because of this.

One trick that works is to use a non-magnetic filler to eat up any space on the ends of layers. Nylon monofilament fishing line comes in diameters of about 0.015” for the 15lb rated lines. This makes a GREAT spacer, being almost exactly the size of the wires that are being left out.

There are also tricks to starting and ending windings. To start a winding fold about an inch of tape over the wire so the wire is held, but there is a sticky-side-down bit left over. Place the taped wire in position on the bobbin, and wind a couple of turns next to the first one and on top of the tape. This holds the first wire down.

Bobbin Directions and Faces

There are only two faces of the bobbin that are usable for bringing out lead wires. It's important to decide before you start winding which face will bring out which lead wire. Decide ahead of time which face will have the primary leads and which will have the secondary leads coming out. It's a lot of work to go change this after the fact. I advise putting the primary leads on one face and the secondaries on the opposite face, because on a small transformer like this, the faces get crowded.

Starting a winding

Here's a quick picture of how to tape down the first wire:

A bit of tape (yellow) is folded over the wire. This held in place by the second and following turns on the bobbin.

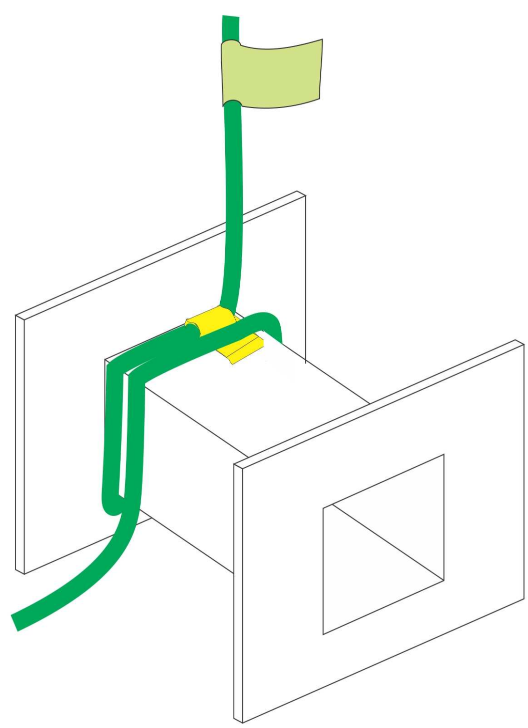

Notice that the wire coming to the outside of the bobbin to attach to a lead has been labeled with a bit of masking tape so you can remember which wire is the start and which is the finish. This will be important when you start putting leads on.

Winding Instructions

The average (mean) length of a turn on this bobbin is about 5", so you can compute how long each winding will be. This is a conservative estimate for all but the outside layers.

1. Set up the #29AWG wire for winding. Tape down the start of the bottom primary section on the bobbin. Tag the start of the primary with masking tape and a pen.
2. Wind a layer of 74 turns, making sure that each wire lies smoothly touching the wire next to it.
3. Add end turns of monofilament fishing line and tape over the first wire layer one layer of tape.
4. Wind on a second layer of 74 turns. Again, use fishing line to fill the ends so the next layer will not spill off the ends.
5. Hold in place with one layer of tape.
6. Wind a third layer of 74 turns; at about 66-68 turns, lay a piece of tap sticky side up on the bobbin and finish the layer. Fold the tape back over the layer to hold the last turn, hold down with a second piece of tape, and lead the end of the first half-primary out of the coil, marking it with masking tape as the end of the first half primary.
7. Put one layer of tape over the half-primary.
8. Put on one layer of mylar sheet.
9. Measure off and cut 72 feet of #27AWG. Wind this neatly on a spool, taking care not to kink it. Parallel the free end of this length with the end of the remaining wire on the spoor of #27. Tape these together with masking tap and mark this tab as "Secondaries start".
10. Secure these parallel wires with tape to the bobbin, and lead them out of the winding area, again securing with tape. Wind on 30 turns of the parallel wires, taking great care to make sure they lie parallel and do not cross over one another or become twisted.
11. At the end of the first secondary layer, wind on one layer of tape, and wind another layer on top of this one, again taking great care to make sure the wires are wound in a level layer side by side and do not cross over or twist.
12. Tape this second layer down, and repeat until you have five layers of 30 paralleled turns. Cover with one layer of tape. Wind on 17 double turns on the sixth and final secondary layer, spacing the turns out so they cover the entire layer evenly. Use fishing line for help in spacing evenly if you need to do so.
13. Cover the final secondary layer with one layer of tape.
14. Put on one layer of mylar sheet.
15. Repeat steps 1-6 to add six more layers of the second half of the primary.
16. When finished, cover the outside of the coil with two layers of tape.
17. Using the masking tape notes, identify the end of the first half primary and the start of the second half-primary. Trim these ends to a convenient length, sand off the enamel, twist them together, and tape this connection down to the coil.

18. Using the masking tape notes, identify the start and end of the secondaries and primaries.
19. Attach leads to the starts and ends of the joined primary and both secondary windings. Secure these to the coil with tape, and put on the final outside layer of tape.

Re-Assembly

When the coils are wound and taped, leads attached, all ready to go, you get to stack those laminations back inside it. Fortunately, we do not want the interleaved stacking that was used in the original donor. In fact, that will not work properly. This transformer needs a specific gap.

So stack the Es back inside the bobbin, with all the legs pointed the same way. Stack it as full as you can with the laminations you have. Where a lamination is not flat, reserve it for last, using all the flat ones first. Clean up any gross flakes of varnish and such so the laminations lie as flat on one another as possible.

[talk about plan B here in case the coils are outside the bobbin edges...]

When the bobbin is as full of laminations as you can reasonably get it, stack the Is in a block. Place the Is in the top of the frame as a block. Insert your gap spacer sheet, then slide the stacked coil and Es into the frame so the ends of the Es' legs touch the gap spacer. Joggle the core laminations as square as you can get them, then place the bottom plate (if any) onto the bottom and tidily fold over the retaining tabs. You can use light taps with a piece of wood or rubber hammer to tighten things up.

Testing

Your transformer is ready to test. Use you ohmmeter and verify that the primary is continuous from end to end and that it does not short to either secondary, nor secondary to secondary. Do this test: http://geofex.com/FX_images/xform_test.gif to verify that there are no shorted turns. If it fails these, rip it apart and rewind.

If all seems well, try it in the amp. If it runs, remove it from the amp, and impregnate it with varnish or wax. If it doesn't, do the debug to find out why. I personally have issues in that there are three different windings that I can get hooked up backwards!

Potting

As noted, there are two processes, wax and varnish

Wax impregnation is an old, old method. You melt the wax in a double boiler and then lower the transformer into it, letting the transformer get completely hot so the wax penetrates it fully. The transformer is ready for use when it cools. The bad thing about wax is that if the transformer gets hot inside, the wax will re-melt and run out. This is a common issue in some AC30 amplifiers from the 60s.

WARNING! Do not try to melt wax in a pan over an open flame without using a double boiler to keep the wax from getting to a high temperature, vaporizing, and starting a flash fire with the wax vapor.

Varnish impregnation is what is used on modern transformers. The process includes heating the transformer in an oven to about 140-160F, then immersing it in the varnish. The heating is to keep the varnish thinned by the heat until it can run all the way inside the coils. When the transformer is good and soaked full, it's removed and allowed to dry a bit, then encouraged to dry by being baked at a low temperature, again perhaps under 200F in an oven.

Varnish is the class way to do it. However, it's difficult to find the right varnish, perhaps expensive, it's messy to muck about with that much drippy varnish, it's smelly to do in a kitchen where there's an oven, and it takes a long time.

Beyond that, if you're doing a super job, you will pull a vacuum on the varnish pot with the transformer inside it so the air comes out of the transformer and the varnish goes in when pressure returns. Modern high class transformers are vacuum impregnated like this.

Getting the right kind of varnish may be tough. There are places where they'll sell you no-fooling electrical varnish. It's about $40 a pint when you can find it. But you can also use “short-oil” alkyd or phenolic varnish. These are not the home center polyurethane varnishes. Woodworking sites sell the right kind, generally as a hard-finish varnish for things like table tops. I intend to try out Behlen “Rockhard” tabletop varnish, which is a short-oil phenolic type.

If you varnish, be sure to bake till you're sure it's dry and cured inside as well as out. Drippy varnish is very, very messy.

Neither of these processes is civilized enough to perform in the average kitchen. The are messy and smelly, and have the potential to start fires unless done carefully. Be careful about where you do your work as well as how.

Appendix A. Just Tell Me The Turns

1. Obtain a donor transformer on 75EI laminations with a nylon bobbin and external frame, and winding supplies.
2. Disassemble the donor into laminations, frame and bobbin.
3. Wind 222 turns of #29AWG in three layers. You may choose whether to use tape or mylar sheet inter-layer insulation to keep a level wind without turns falling off the ends.
4. Wind 167 turns bifilar of #27AWG in six layers. Make the last layer of 17 turns spaced over the entire length of the winding area.
5. Wind 222 turns of #29AWG in three layers.
6. Connect the finish of the first winding to the start of the last winding and tape inside the core to form one split winding of 444 turns.
7. Attach and secure six leads, one each to start and finish of all the remaining wires.
8. Finish the coils with external tape.
9. Insert laminations into core, butt stacking the Es, and using a 0.007” paper spacer between Es and Is.
10. Insert coil and core into the frame and secure the frame tabs.
11. Test for function and correct if needed.
12. Upon successful testing, pot in wax or varnish if you want.

Appendix B. Thomas Organ Vox Totem Pole Circuit

Here's a redrawn schematic of the V1141 power amplifier circuit.

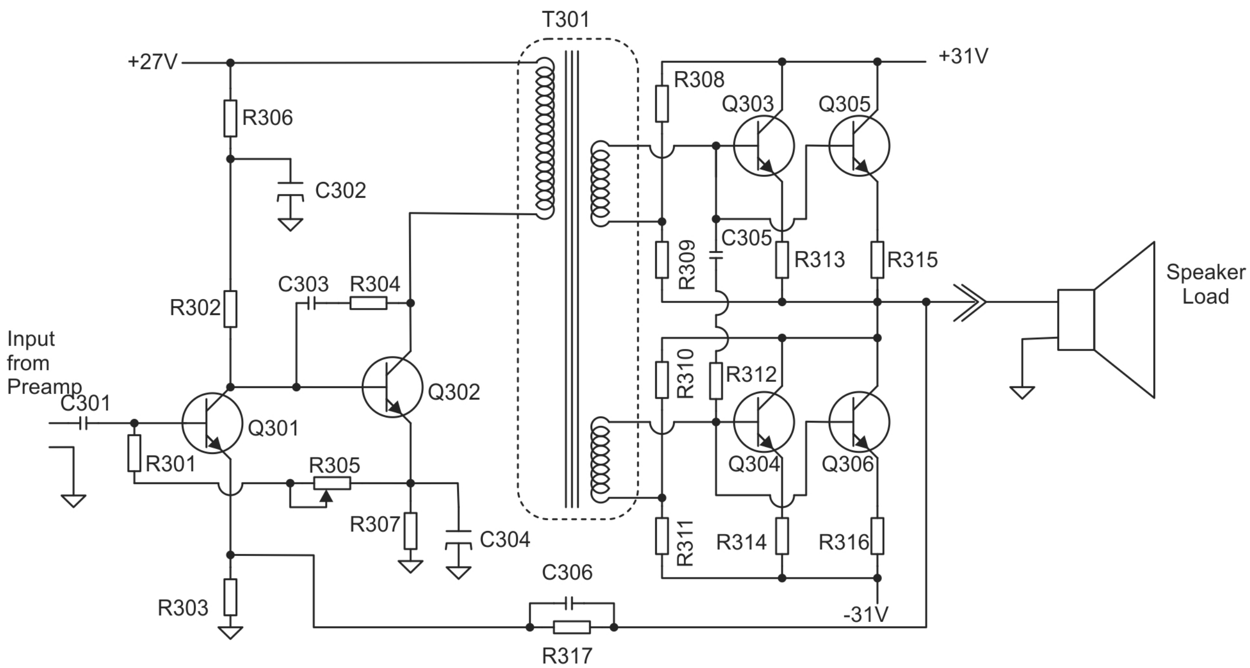

Schematic Per Service Literature

How does this thing work? Here's a simplified version for some analysis. I removed one pair of power transistors to make it like the Royal Guardsman.

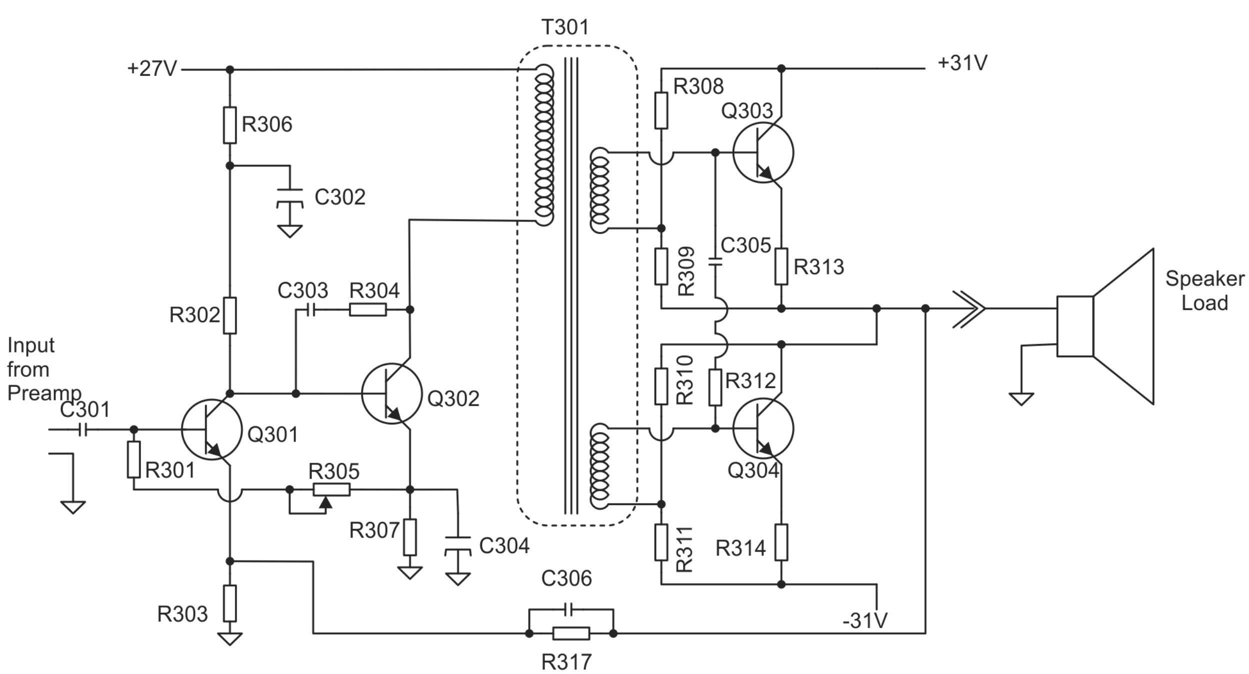

Simplified for analysis (Royal Guardsman V1131 is similar)

The output side is a stack of two of the circuits shown at right. The power transistor is biased across the 31V or so of the power supply by a voltage divider of R308 ad R309. If the resistors, transformer windings and transistors are identical, then a stack of two of these will have a junction at the speaker load output which sits exactly at half the total power supply voltage. Since one is on the positive supply and one on the negative, the output sits at nearly 0V. It will be off by only the amount of the non-identical-ness of the two sections.

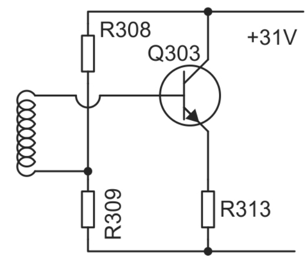

Bias Analysis

Q303 is biased slightly into conduction by the voltage across R308 and R309. The other output transistors are similarly biased. Under nominal conditions, the voltage at the output will be 0V. So the voltage across R309 is Vbias = 31* R309/(R308+R309).

This voltage flows through the resistance of the transformer winding, the base-emitter junction of Q303, and then through R313. For the design values of R308, R308, and R313, the voltage across R309 is Vbias = 0.54V. Under normal conditions, about 30-40mV appears across R313, so the actual bias impressed on Q303's base-emitter is about 500mV.

R308 is much larger than R309 (220 versus 3.9 ohms), so the value of R309 has hardly any effect on the current that flows through the bias string. Accordingly, one can adjust the transistor current by raising or lowering the value of R309.

Stacking two of these across a balanced bipolar supply puts the junction of the two – which is the speaker output – at the same DC voltage as the ground for the two supplies. If one of the two has off-tolerance resistors or higher gain transistors, the speaker output will be higher or lower than ground.

This can be tweaked by changing R309 or its alternate in the negative side, R311. It is simpler to parallel one of these resistors than to add more in series. Paralleling R309 or R311 with another resistor lowers its effective value, lowers the bias voltage on that transistor, and so lowers the current that transistor allows through. Paralleling R309 lowers the speaker output DC voltage. Paralleling R311 raises it.

Never parallel both of them. This cuts down the total idle current in the stacked stages, and can easily introduce crossover distortion.

The current through R308-R309 and R310-R311 is related to the total current the transistor conducts at idle. So in the Beatle, with two transistors running at about 160ma each, the current through R308-R309 is 31V/223.9 = 138ma. In the Royal Guardsman, there are only two output transistors to be biased, and so only half as much bias string current is needed. Accordingly, the bias resistors analogous to R308-R309 and R310-R311 are about twice the resistance.

Drive Current Analysis

The output of a single transistor pair drives a 4 ohm load in the Royal Guardsman. The Beatle simply adds another output pair in parallel. For 60Wrms into four ohms, the voltage must be about 15.5Vrms. The current is about 3.9Arms, or 5.47A peak. I have measured the output of a Beatle and Royal Guardsman, and they can (sometimes) get the output into a resistive load to +/-28V peak. For a 4 ohm load on a single pair, that would be 7A peak. I'll use

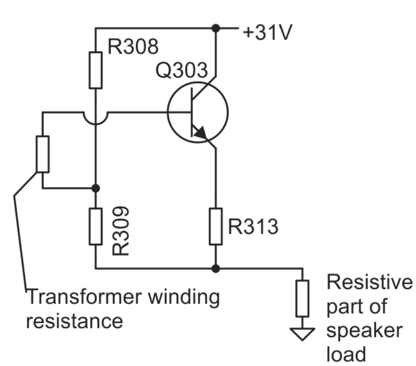

7A as a design point.

Q303 was a selected version of the 2N3055. Squinting mightily at the “typical” curves in the On Semi datasheet, the 3055 has a gain at 7A of about 20, so we need to provide 7A/20 = 350ma of base drive. The transformer winding has to provide a voltage that will drive this current.

The signal voltage must be at least as large as the voltage across R313, Vbe for Q303, and the base current flowing through the transformer winding resistance. With 7A flowing through it, R313’s voltage is 7*0.25 = 1.75V. Vbe from the datasheets at Ic =7A is about 1.2V. So the base of Q303 must be 1.75+1.2 = 2.95V above the output voltage, and 350ma must flow through it.

That same base current must flow through the transformer winding and the parallel combination of R308 and R309. We can approximate this by just R309 with little loss of accuracy. I measured the Beatle transformer winding at 6 ohms. So we need to provide voltage to force 0.35A through 9.9 ohms, or about 3.65V, in addition to the 2.95V for the base-emitter and R313. This gets us to a peak voltage on the secondary of about 6.42V.

The secondary is then providing 350ma peak at 6.42V peak. That’s a sine-wave equivalent of 4.54Vrms at 0.247Arms, or similar to the current and voltage into a load of 4.54/0.25 = 18 ohms. The Beatle driver transformer does much the same, but that would make for a similar load of 9 ohms. Hence my comment that this is a small audio output transformer in its own right, putting out over 1W for one pair, 2W for two pairs.

The Beatle driver transformer has to do 350ma per transistor, so it must supply 700ma of base drive.

On the primary side, the primary voltage swing must be at least the turns ratio times the secondary voltage. If the secondaries have peak voltages of 6.42V, then the primary peak voltage must be a bit over 17V. The primary is driven from 27V, and loses 3.5V to the bias resistor at the emitter of Q302. This leaves about 22V to get 17V for the primary and any losses in Q302.

Since this is a single ended stage, the output swings upwards an amount equal to its swing downwards, so the primary can swing +/-17V around the +27V power supply for the transformer primary.

The primary bias current is 3.5V across 10 ohms, or 350ma. There is a 2.67 current step up to each secondary, so each secondary can source .35*2.67 = 0.934A into the bases of the output transistors. From our earlier analysis, this is enough for each of the Beatle transistors to be fully driven.

It occurs to me that one failure mode in the driver is when an output transistor base goes open. The flyback voltage on the primary could go to quite high voltages, possibly killing the driver. This doesn't happen often, but I think it could. There is no protection for this in the design. In fact, there is no active protection in the designs at all.

These designs were OK at their time, but with the perspective of today's devices and techniques, it's possible to upgrade the designs – a lot! I have some ideas about that, but that's for another time.

Appendix C. Detective Work On The Beatle Transformer

It is possible to measure a lot of the characteristics of a transformer without actually seeing inside it.

I have both a Royal Guardsman and a Beatle to play with, but I've so far been too wimpy to destroy one of them to find out what's inside the driver transformer. So I tried more indirect measurements.

First, I measured the resistance of the primary and secondary windings.

Next, I unsoldered all the leads and applied an AC signal voltage to the windings. A transformer causes a voltage on every winding proportional to the number of turns on each winding. If one winding has, for instance, 0.5V appearing across each turn, then all turns will have that same voltage, in the ideal case. The ideal case is approximated by there being no loading on any winding, and by the frequency of the signal being in the middle of the transformer's frequency range, far away from the effects of low and high frequency cutoffs.

I managed to thread 40 turns of wire-wrap wire into the spaces around the bobbin on the driver transformer I was measuring. This let me put a signal onto the primary, then measure the voltages on the voltage on the 40 turns, and then measure the voltages on the other windings. Once I know the voltages, I know the turns ratios between the windings, and knowing the voltage across 40 turns lets me know the volts per turn. I can then compute the number of turns in the primary and two secondaries. The info is subject to how accurately you can measure the volts on each winding. There is still some uncertainty based on the measurement errors. But you can get very close.

The primary winding was driven with 4.0V peak (8 pk-pk) of 1kHz. I set that to be 4.0 as close as I could with the signal generator level control. On the other windings I measured:
Orange-yellow, 1.5V peak; same for white-gray. The red, orange, and white leads are in phase; they all go up and down together in synchronism. The 40t sense winding had a voltage of 0.38V peak, (0.76V pk-pk).

The voltage ratio from sense winding to red/black primary is then 4.0/0.38 = 10.52:1. That makes the primary turns 10.52*40 = 421 subject to any measurement errors. The two secondaries are then (1.5V/0.38V)840T = 157 turns each.

Based on some analytical work I did, I'd start with a target of 420 turns for the primary, and 158 or 159 turns for the secondary.

Physical measurements of the transformer and stack give a lot of the physical details such as the lamination size and stack.

Appendix D. Sizing Wire and Insulation To Window Area

Getting enough turns of the right (or big enough) wire to fit in the area available inside the window of a lamination is somewhat tedious. There are many possible answers. Picking a better if not best one can be a lot of iterative work.

75EI has a winding window of 1.125" by 0.375". The bobbin eats up some of that, so you can measure inside your bobbin for some estimates. My bobbin measured 1.025” wide by 0.306” tall where the wires go.

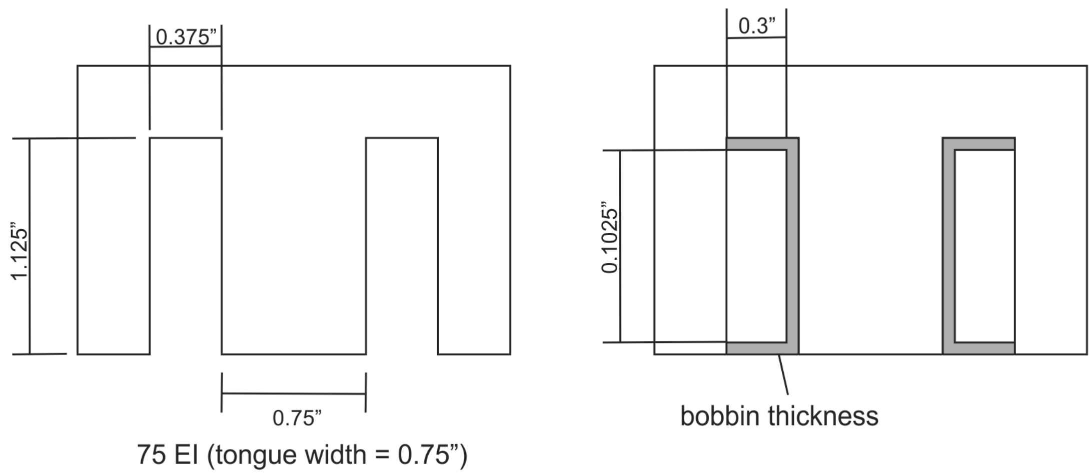

One layer of wire wound side by side can only be the bobbin length (1.025”) divided by the wire diameter, including insulation. There are detailed tables of wire properties available. For instance, if we want to know how many turns of #25AWG will fit in one layer on this bobbin, we look up the diameter of #25 and find that there are several diameters listed. There are single build, heavy build, triple build, and a minimum, average and maximum on each of these. It is a truism that you can never wind as tightly and compactly as you think you can, so I used the diameter for maximum heavy build in all my calculations.

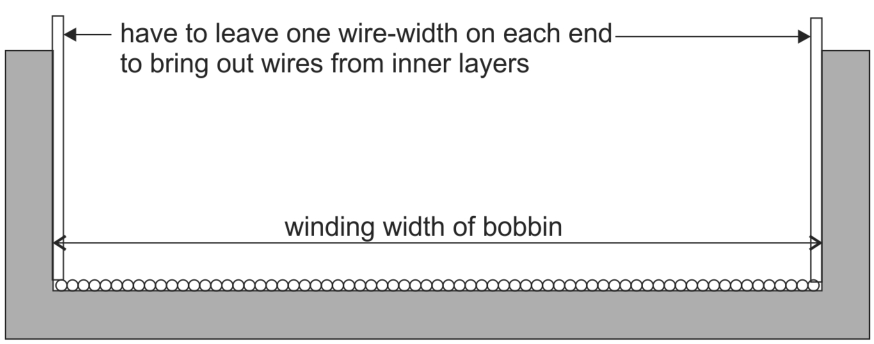

So on one layer, 1.025” wide, with 0.0203” diameter #25, we can get N = 1.025/0.0203 = 50.49 turns. Actually we can only get fifty turns before the wire starts climbing on top of the previously wound turns. So we can get 50 turns per layer of #25.

If we're trying to wind a primary with at least 420 turns, that's 420 turns divided by 50 turns, or 8.4 layers. This means that there will be eight full layers and one only 4/10 full. It will build up to 9 layers high, though.

For magnetic-field reasons, it is generally bad to have partial layers, and when you do have them, it's best to space them out over the full winding width of the layer. So if we used #25 for the primary on this transformer, it would be 18 layers tall for just the primary. That's 18 layers times 0.0203" or 0.3654" which is taller than the available space. It won't fit, even just the primary, if we try to use #25.

Those are the basic calculations. You pick a wire size, look up its diameter, compute how many turns you can fit on a layer, then how many of those layers get you the number of turns you need. When that is done, you compute the height that this many layers of wire (plus any inter-layer insulation you use) will build to. That has to be less than the height available in the window and bobbin. It is smart to leave yourself some space for tape and additional insulation, because perfect winding is nearly impossible.

Most magnetics designers start with the available winding window area, divide it up proportionately to the power in each winding, diminish that by some safety factor based on experience, and then compute a starting wire size from that.

In this bobbin, there is an area available for wire of 0.3825 square inches. We'll start by allowing 0.19125 square inches for the primary and the same for the secondary. We'll also drop off some of those decimal points when we can.

We need to fit 420 turns into 0.19125 square inches. It would be smart to cut that to about 85% of the total, or about 0.1626 square inches for a starting point. If this amount of space holds 420 turns, then each turn has 0.1625/420 = 0.000387 square inches. Taking the square root, we get the biggest wire diameter that can possibly do that. That number is 0.01967". The next smaller wire size is #26, at 0.0182" (maximum diameter, heavy build).

With that to start, we calculate turns per layer, layers, and height build. The height build had better be less than half the available 0.0306" height in the bobbin, or this will will not fit either. And we'd better leave room for at least some insulation between layers to keep the wires laying in flat layers. This quickly gets tedious on a manual calculator. You can imagine how tedious it was on a slide rule.

I did a quick spreadsheet that when given the winding length, wire diameter, and available height, computes the number of turns per layer, layers, and build height.

Actually, I copied the formulas so I could do this once for the primary and once for the secondary, and the spreadsheet would compute the total turns, layers, and build height. With this mechanized, I was free to play with several what-if situations.

So here's where I started:

Primary Turns	Primary AWG	Pri T/layer	Pri Layers	Pri Height	Sec Turns	Sec AWG	Sec T/layer	Sec Layers	Sec Height	Total Height
420	26	54	7.8	0.1696	316	26	54	5.9	0.1272	0.297

Wow – a fit on the first try. But notice that the total height is vanishingly close to our maximum height, and there is no room left over for extra turns of tape and such. It works in theory, but it's not windable by a beginner. Let's try a smaller primary size.

Primary Turns	Primary AWG	Pri T/layer	Pri Layers	Pri Height	Sec Turns	Sec AWG	Sec T/layer	Sec Layers	Sec Height	Total Height

420	27	60	7.0	0.1365	316	26	54	5.9	0.1272	0.2637

Better. But there's only 0.033” of height left over. Still not easy to wind. Next try.

Primary Turns	Primary AWG	Pri T/layer	Pri Layers	Pri Height	Sec Turns	Sec AWG	Sec T/layer	Sec Layers	Sec Height	Total Height
420	27	60	7.0	0.1365	316	27	60	5.27	0.1170	0.254

Also better, as now there is about 0.05” for outer tape to hold it together.
But we notice that although the primary is seven layers, the secondary is just a bit over five layers. The primary doesn't have to carry as much current as the secondaries. Let's go smaller on the primary to leave more room for the winding.

Primary Turns	Primary AWG	Pri T/layer	Pri Layers	Pri Height	Sec Turns	Sec AWG	Sec T/layer	Sec Layers	Sec Height	Total Height
420	28	67	6.3	0.1239	316	27	60	5.27	0.1062	0.2409

With the primary smaller, we could either go back to #26 for six layers, or down to #28 for five layers. Going up to #26 gives six layers for the secondary, and a total height of 0.2511.Going down to #28 gives five layers (4.7 actually) of 67 turns per layer.
But there are other things we can do. Transformers generally need enough turns, not the specific number. Since we've already spent seven layers high on the primary, what happens if we make it seven layers even of 67 turns per layer?

Primary Turns	Primary AWG	Pri T/layer	Pri Layers	Pri Height	Sec Turns	Sec AWG	Sec T/layer	Sec Layers	Sec Height	Total Height
469	28	67	7.0	0.1239	352	27	60	5.87	0.1170	0.241

This requires us to increase the secondary turns to 352 to keep the same ratio. And we get seven even primary layers, nearly six secondary layers, and a height of only 0.24”. This is pretty good. It's a usable alternate that suffers only from having an odd number of primary layers. If we were trying to do a no-expense-spared design, we could split the secondaries into two sets of three layers, and the primary into three sets of layers, two, three, and two. That's probably getting too fancy, but it's possible.

Back in the real world, we have extra primary turns. What happens if we want only six layers on the primary? Six layers at #28 is only 402 turns, and going down in turns is not good. Let's try #29 and 420 turns.

Primary Turns	Primary AWG	Pri T/layer	Pri Layers	Pri Height	Sec Turns	Sec AWG	Sec T/layer	Sec Layers	Sec Height	Total Height
420	29	74	5.7	0.0978	316	27	60	5.27	0.1170	0.215

Hmmm. 74 turns per layer, total height of 0.215”, easier to wind and somewhat forgiving. Better.
What if we used six full layers of #29 for the primary?

Primary Turns	Primary AWG	Pri T/layer	Pri Layers	Pri Height	Sec Turns	Sec AWG	Sec T/layer	Sec Layers	Sec Height	Total Height
444	29	74	6.0	0.0978	334	27	60	5.57	0.1170	0.215

OK, we get more than the minimum turns, an even number of primary layers, a nearly-even six layers of secondary, and a total build height that's not a challenge to wind.

I stopped there. That's where I'll wind mine, and that's what I recommend in the write up.

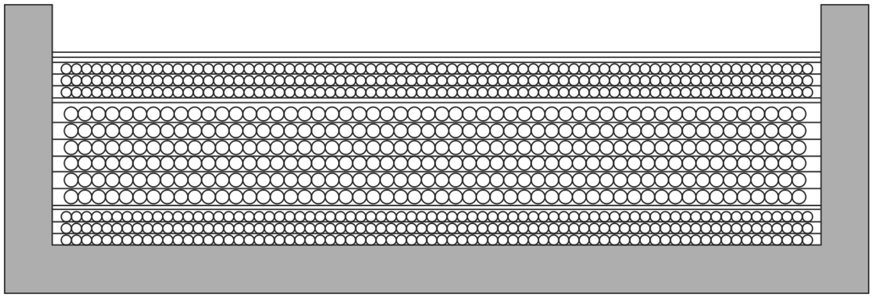

This is what we want the winding to turn out like. In reality, it's never that easy.

Here's one final possibility. If you change back to #26 for the secondaries, and change the number of primary turns to 432, you can get to

Primary Turns	Primary AWG	Pri T/layer	Pri Layers	Pri Height	Sec Turns	Sec AWG	Sec T/layer	Sec Layers	Sec Height	Total Height
444	29	74	5.8	0.0978	324	26	54	6	0.1272	0.225

Which is also a usable wind, and maybe slightly better as it gets more copper into the higher current secondaries.

As a note, I've been computing what I call “usable turns per layer”, which is the total turns per layer computed, rounded down to an integer, then with two turns subtracted for space to get wire out on each end.