So You Thought Your Amplifier Was Balanced?

No matter what the advertising or manual says it probably isn't!

Foreword

Balanced connection has become one of the fashion words in the audiophile vocabulary, however, no matter who one asks, no-one seem to really know what it means, so Andy Grove and I thought it was on time someone explained what the term balanced line actually means, why it is used and what constitutes a proper balanced connection, let's first start with some basics.

What is the purpose of using a balanced interconnection in audio?

Its original purpose was to combat the effects of hum (and noise) pickup in long cables, for example those used in recording or early TV studios where cables can be very long. Hum is a form of EMI, Electro Magnetic Interference. Nowadays the aether is jam packed full of it, from low frequencies, such as the power line frequencies of 50/60Hz, right up into the microwave region, which goes to and beyond the gigahertz region (1Ghz=1000,000,000Hz).

A bit of boring theory is necessary, but to get back to physics style basics:

1) When we talk about electricity we are talking about charge with electrons as the carriers.

2) Static (not moving) charge creates a static Electric (E) field.

3) Uniformly moving charge creates a static Magnetic (H) field.

4) The non-uniform movement of charge (in other words AC) creates an electromagnetic wave composed of varying E and H fields. The fields are connected in a vector fashion and are at 90 degrees to each other.

At mains frequencies, generally 50 or 60Hz, the wavelength of the electromagnetic wave is enormous, around 6000 km. Normally we are interested in noise picked up at distances very considerably smaller than that, and by electrical components, which are physically a tiny fraction of that wavelength. At these distances the electric and magnetic field vectors of the electromagnetic wave are experienced separately.

The varying (at 50/60Hz) electric field induces a voltage directly onto things, which are essentially insulated from ground (in which no or only a small current can flow to ground), this happens by CAPACITIVE coupling. An example would be our bodies, and it's that which you hear when you touch the high impedance input of an amp and get that BRRRP! All the wiring and so on around us acts as one plate, and our body as the other plate in a capacitor. The coupling isn't great and so the effective capacitance is tiny, depending on the area of the "plates" and the distance between them (which is large), so when you touch a conductor which is "grounded" a tiny current flows. However, my brother regularly climbs trees near super grid power lines he's a tree surgeon, and because the AC voltage on those power lines is very high (hundreds of kilovolts) he will sometimes get a jolt off a ladder or something else which is grounded, in his unfortunate case the current is no longer so tiny.

It's the varying magnetic field, which creates most of the hum in our cables.

The effect is similar to that of transformer action, the primary of the transformer is all the mains cabling, motors, power transformers and so on, and the secondary of the transformer is the wire in our cables. This is the INDUCTIVE coupling. The effectiveness of inductive coupling is once again dependant upon distance, this time area is important but in a different way. If we are talking about a cable, then the area enclosed by the loop of the two conductors (and of course, by definition, their length), which make up the cable is the relevant factor.

At higher frequencies, radio frequencies, the wavelength is very much shorter, the electric and magnetic vectors of the wave are combined, and we see a true Electro Magnetic wave at normal “human” distances. EM waves are propagated according to Maxwell’s equations.

Ever wondered what the screen on your cable is for? To screen out hum?

Well basically it doesn’t do much at all, not to hum anyway. An electric field can’t penetrate a conductor, so the screen is effective for the E field part of the overall hum field, and it’s also good for eliminating radio frequency noise. But, in reality, the area of the two conductors in a cable is small compared to their length, so magnetic coupling is the most important problem here, and the H (magnetic) field goes straight through that thin screen. A static or DC magnetic field is unaffected by normal conductive materials, and AC fields are only affected due to the induction of eddy currents in the material.

So how do we combat the magnetic component of the hum field?

We could screen the cable with a material, which was a “conductor” of the magnetic field, materials which have high permeability (or magnetic conductance) are generally ferrous (containing iron) materials. The highest practical permeability occurs in the nickel iron alloys containing around 80% nickel, Mumetal is an example. If air is taken to have a relative permeability of 1, Mumetal has a relative permeability in the hundreds of thousands, up to maybe 500,000. Some nanocrystalline or amorphous materials may be able to claim even higher raw material permeability, up to 1 million, but they are extremely brittle (the term metallic glass is often used) and the raw material figure is never achieved in practice. The relative permeability of copper. aluminium and silver, the commonly used conductive metals is of the same order as that of air, 1, so their effectiveness as a magnetic screen is severely limited, unless that screen is really, really THICK.

It is very common to use a Mumetal screen to shield small microphone or moving coil cartridge step-up transformers, and even for special applications to wrap a cable in a Mumetal screening tape, the problems with this are that Mumetal is very expensive, adds considerable bulk because it has to be fairly thick and thus makes the cable heavy and inflexible, and because the cable is surrounded by a magnetic material its electrical characteristics are altered, sometimes unfavourably, at least when we are talking about transmitting audio signals.

Usually we have two conductors in our interconnect cable, to make the circuit, send and return, or signal and ground. At one end you have your preamp, which will have a reasonably low output impedance figure, usually much lower than the power amp input impedance, so for simplicity we can say that at that end of the cable the two ends are joined together. What you power amplifier, at the other end of the wires, does is measure the voltage between the ends of the wires, and amplifies it. The wires are effectively connected together at the preamp end and it’s not too difficult to imagine that you have a loop, from the power amp signal input, to the preamp, through the preamp output stage and back to the power amp input ground. This loop will work quite nicely as a turn in a transformer with all the mains cabling and so on as the primary. If you pull the two wires further apart, you make the area enclosed by the loop larger and the voltage induced between the two ends at the power amp will be larger, and you get more hum. Put the wires as close together as possible and you minimise hum. The way to get them as close as possible together is to twist them together, this, strangely enough, is called a “twisted pair”.

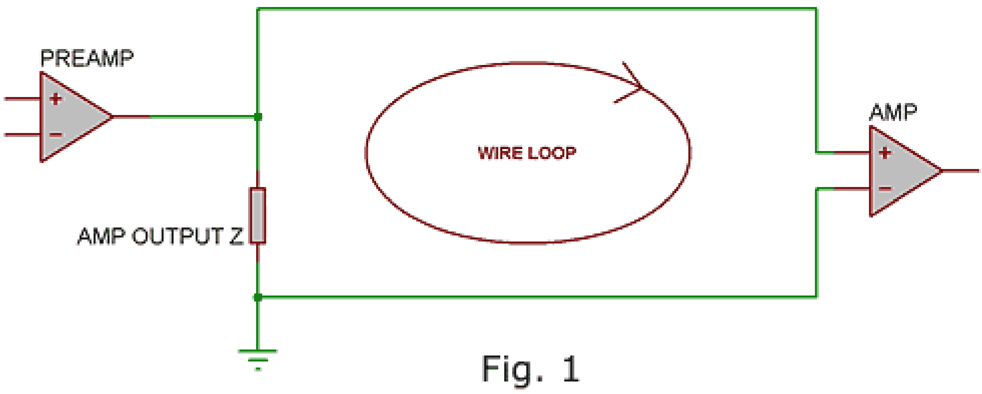

Fig. 1

Unfortunately for most domestic audio equipment there is another loop to consider, and quite often it's this loop, which causes all the problems. Usually both preamp and power amp will be connected to ground by their respective mains power connections, and unless the unit is double insulated this is a legal requirement in most countries. Now we have a wire from each unit going down to ground and another low impedance connection between the two units, the signal ground of the cable. This forms a loop a big one and a current will flow around it, again like the secondary of a transformer. This current will, due to the resistance of the conductors and connections, generate a voltage at the power amp input, and you get hum, this is the dreaded "ground loop". Of course the situation is magnified at a phono stage or microphone input because the effective gain is so much higher. Dirty or corroded connectors at audio interfaces only make matters worse, the current will have to flow through a high resistance layer of oxide, that will generate a higher voltage and bingo: You have HUM!

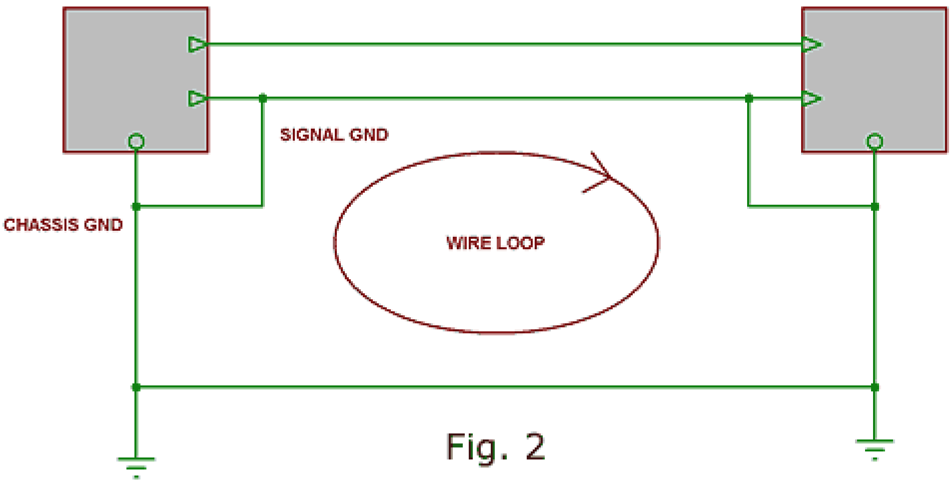

Fig. 2

To recap:

1) Hum in cables is predominantly generated by the mains frequency magnetic field, which is all around us.
2) Hum is picked up by loops with big areas.
3) Ground loops between units, makes even more HUM!

So, to minimise hum we use twisted pairs in the Audio Note™ interconnect cables, or at least the two wires are very close together, and to cut ground loops we sometimes remove the ground connection from one unit, which is really quite dodgy, or, which is quite common, the chassis is taken to ground directly, and the signal circuitry is connected to ground via a high power 10 ohm resistor. This resistor breaks the ground loop, and a lot, but not all, of the hum voltage appears across the resistor instead of appearing in our signal path.

When cables become very long, or when very low hum is required these simple measures just aren't enough, there are secondary effects which means that there's still a fair bit of hum interference.

Using the preamp/power amp example again, the power amp measures the voltage between ground and its signal input pin, in a standard single ended system. Of course the ground wire in the cable is connected between the two equipment grounds, and the signal is connected between the preamp output and the power amp input. The problem is that were looking at two separate “environments” for the two conductors, one is high impedance, essentially the amp input impedance, the other is very low, ground. That means that the effective voltage, which appears across the two wires will be different. As the power amp measures the difference between the two voltages, this ends up as hum.

How?

It’s not normally realised, but all amplifiers (and similar devices) have differential inputs, there’s no such thing as an absolute voltage. When we say “voltage” we mean “potential difference”, the difference in voltage between two designated points. In a single ended or unbalanced input one of the inputs is tied to ground, the other is “the input”. However, when the term “differential input” is used we normally mean one where there’s two input terminals, neither of which is effectively grounded, and it’s between these inputs that the signal is measured. A differential amplifier of any sort measures the voltage between two points, you get no output if those points are both at 0V or, theoretically, if they are both at 20V, or 10,000V, or -3,000,000V, as long as they are both at the same voltage.

If the voltage is different you get an output.

The “unbalanced” or single ended input is definitely unbalanced when it comes to impedance match between the two inputs, as stated before one lives in a high impedance world, the other in a low impedance world. This means that the effective induced AC voltage from the ambient hum field is different for the signal and ground wires, therefore there is a difference in voltage between the ends of the wires at the amp end and, guess what, you get hum.

When the term “balanced connection” is used it strictly means impedance balanced.

The two inputs between which the signal is measured have the same impedance when referenced to ground. This is so that the hum signal which is picked up in the two conductors of the cable is EXACTLY the same, if the induced hum voltage is the same and it’s applied to a differential input then you get no output, hum is terminated.

There are still some potential problems, and they generally arise from a misunderstanding of the subtleties involved, or from a lack of understanding of what a balanced connection actually is. What it ISN’T is a symmetrical signal voltage referenced to a central ground point, that would be push-pull, but this is exactly what many consumers, and even engineers, think a balanced connection is.

Looking at our differential input, we assume that if the two input voltages are the same then you get no output, unfortunately this situation only exists in the Platonic world of the mind, Perfect symmetry does not exist in nature, and this situation is exacerbated by “symmetrical” and “balanced” inputs which aren’t even symmetrical.

A typical and illustrative case is the single op-amp type (operational amplifier) “balanced” input circuit. It consists of one op-amp, and four resistors, it’s very commonly used as a differential amplifier, but one biggest bugs is that the equations used to design it usually assume ideal conditions. A zero drive impedance on both inputs, an amplifier of infinite gain and infinite bandwidth and infinite input impedance on both inverting and non-inverting inputs, and components which are matched perfectly as well. The diagram (?) illustrates the case with all the resistors equal in value. R1 and R2 are the gain setting resistors normally used in an inverting amplifier circuit, and R3 and R4 are used in the op-amp non-inverting input to make the common mode input impedance balanced, even though the differential mode input impedances are different on the hot and cold inputs. Altering the resistor values to bring the differential input impedances in line reall screws things up, as the common mode input impedance is thrown out of balance an example is given with R14 modified to give balanced differential mode impedances.

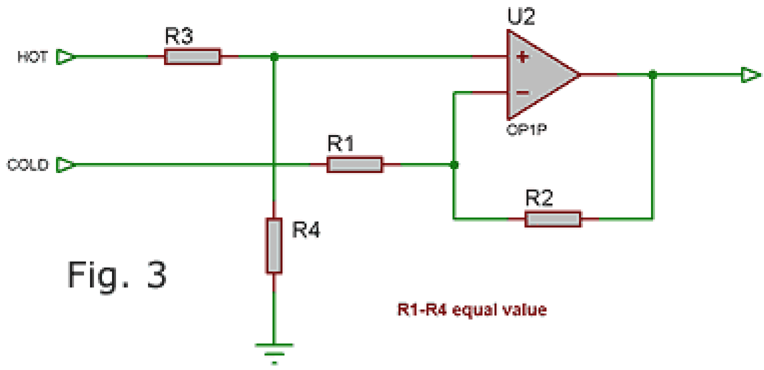

Fig. 3

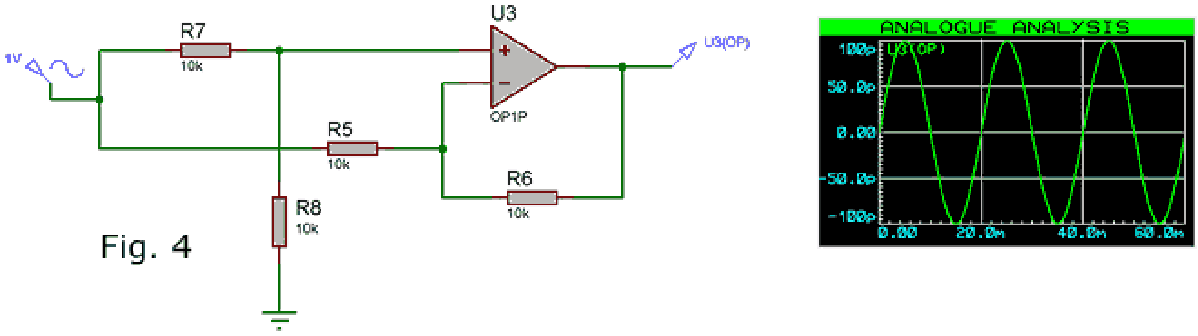

Fig. 4

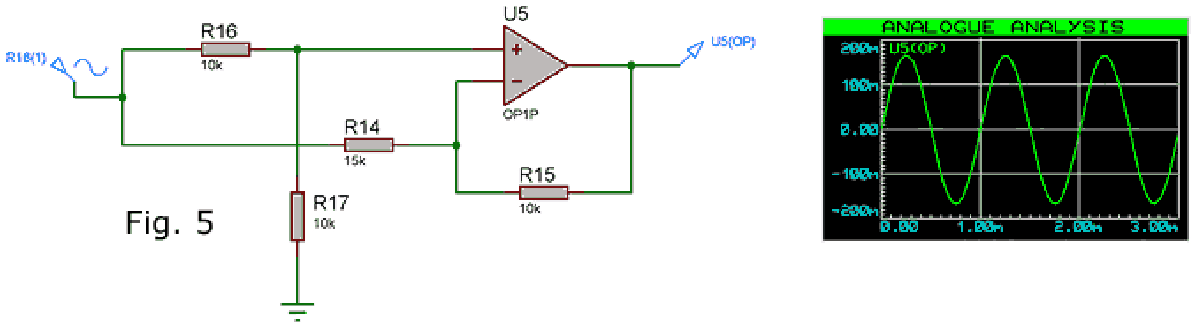

Fig. 5

Conducting the thought experiment on diagram (?) tying both inputs together and applying 1V to both you would get 0.5V at the op-amp non-inverting input, what is the output voltage required to make the inverting input the same voltage? Well, it has to be 0V to get the same division ratio, so with ideal components and the inputs tied together the Common Mode Rejection Ratio (CMRR) is good. The Spice simulation shows an output of 100pV, which is due mainly to the finite gain of the op-amp model (I set it to120dB), in theory it's zero. The circuit is extremely sensitive to tolerances in components, and remember that the output impedance of the line driver and all the connectors and cabling inline is are also components of the system. Let's say all the components were perfectly matched, but one resistor was out by 1%, 1% of 10k is 100 ohms. So I've added a 100R resistor in series with one input. Now the simulation shows an output of 5mV, the input signal used is 1V peak, therefore the CMRR is only 46dB, which is absolutely awful. It's a similar case for the commonly used input (and output) decoupling capacitors as diagram (?) show. Even if the components of the circuit were selected and trimmed it doesn't mean that the guy who made the unit at the other end of the cable did the same thing, and poor connections and so on will degrade the performance of the unit in the same way.

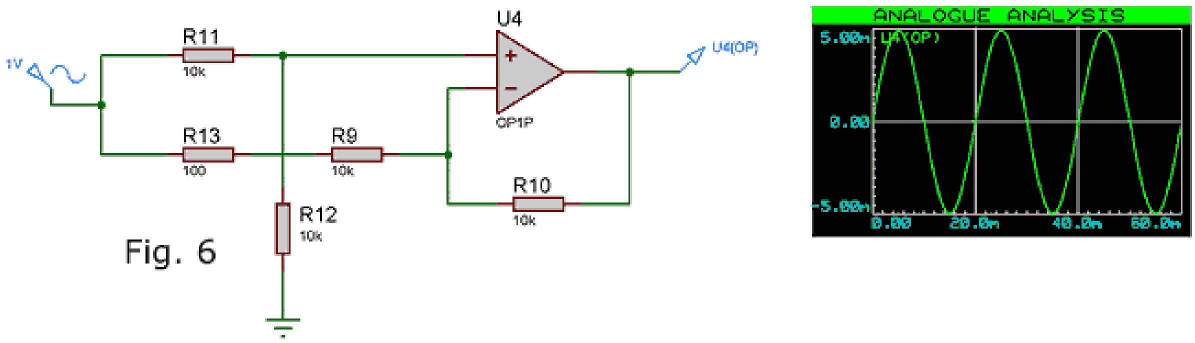

Fig. 6

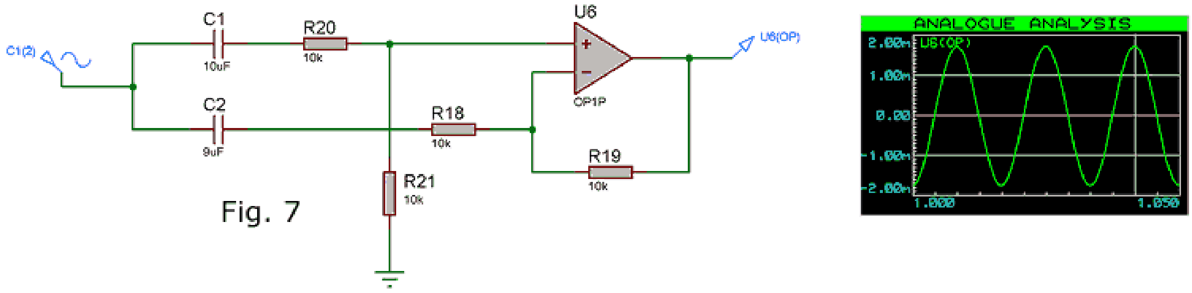

Fig. 7

The problem with this circuit, and all others like it "actively balanced" and so on, is that they are just ground referenced symmetrical inputs, and therefore not balanced at all.

The question we have to ask ourselves is "what happens if we apply a signal to only ONE of the signal inputs and leave the other to float? This is of course the extreme case of what happens in a real world system. And it illustrates the difference between a true balanced system and a pseudo or actively balanced system.

If you look at the model, well, you can see what happens here, and it's not pretty at all. Unfortunately it's the case with just about all of these circuits, it might be an extreme test, but it illustrates the point. The gain on one input is 1 and on the other it's 0.5, so really there's no rejection at all. And it's because in actuality the two inputs are referenced to ground, and they are just two independent inputs connected to form a differential amplifier.

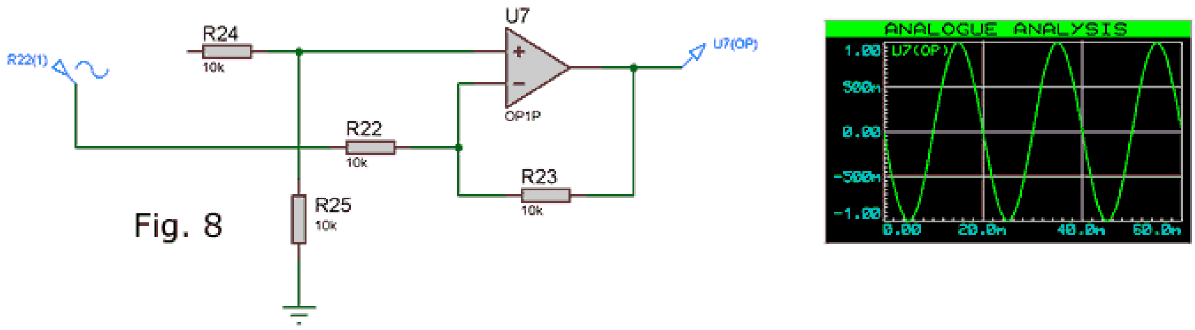

Fig. 8

If you apply a signal to only ONE of the connections of a floating balanced input then still you would get zero output, because the impedance referenced to ground is infinite. However, in practice this is very difficult to achieve, and of course is just about impossible without a transformer input. However with a transformer input, using a well-designed transformer, at 50Hz/60Hz the rejection with unbalanced drive is still very high. The impedance to ground of the open leg is the resistance of the insulation (which will be vast), and the capacitance to ground. It's this capacitance, which degrades the performance of the transformer as the frequency rises. However, the main purpose of a balanced connection is to reject mains frequency noise or hum.

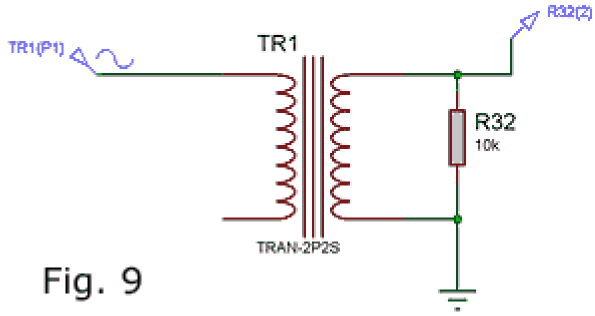

Fig. 9

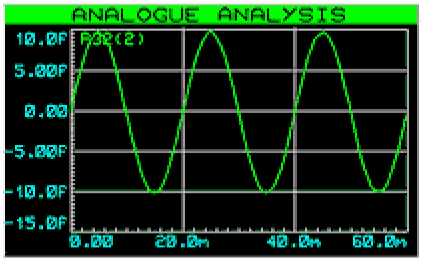

If you look at the drawing with the transformer primary it's easier to see what happens. Perform a thought experiment, taking the limit of a DC signal. If you apply DC to one end of a transformer winding, with the other floating, then both ends will be at the same voltage, and of course no current flows in the winding and you get no output. If the winding had zero capacitance to ground then the situation would be the same at audio frequencies, what you apply to one end appears at the other and you get zero output. However, parasitic capacitance, one of the devil's most devious works, never allows this perfect situation, but it is possible to design a transformer which has a very small capacitance to ground, and at the frequencies of interest the rejection can be made very large indeed, for example, the Audio Note™ M10 input transformer has a rejection ratio of 126dB at 50Hz that's (2,000,000:1) with a balanced source impedance, and 100dB with a serious imbalance of 20 ohms in one leg and 700 ohms in the other.

The model illustrates what an "actively balanced" input does when you do the same to it. Here the rejection ratio is only around 30dB, that's 70dB or 3160 times worse than the transformer. Contrast this to the, albeit more or less theoretically perfect transformer, diagram (?) illustrates the situation with just one leg driven, the only output which is there is due to the requirements that Spice has about large resistances, it can't deal with infinity obviously.

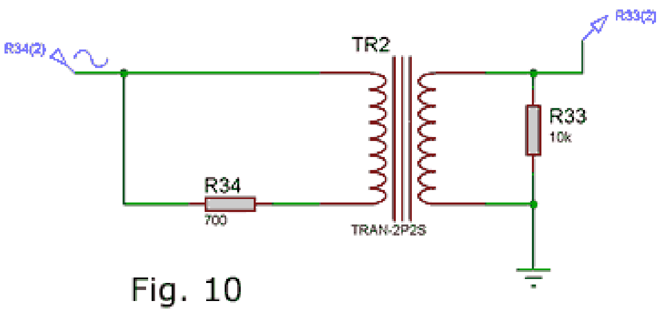

Fig. 10

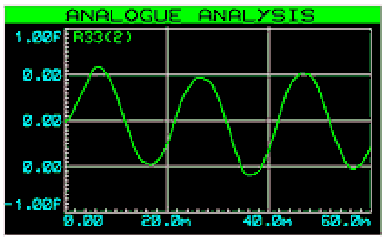

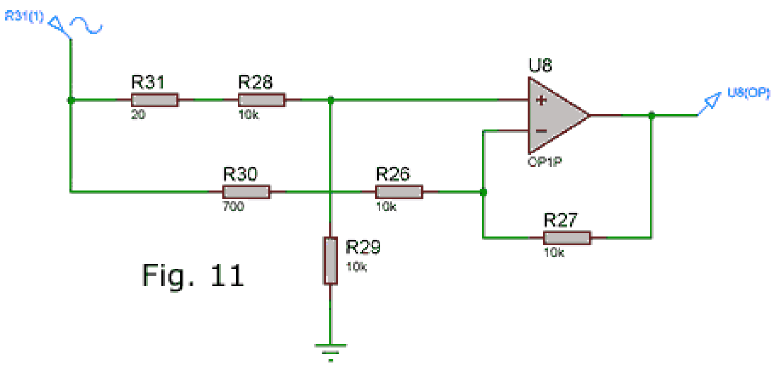

Fig. 11

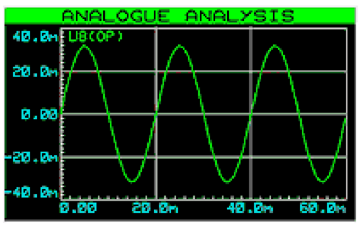

So the majority of these “balanced inputs” are in fact just two single ended inputs, there is no “balanced” about it.

If you can imagine two equal weights, when they sit on the butcher’s counter are they balanced?

No.

If you pick one up does it influence the other?

No.

It’s only when the weights are placed one on each side of the scale are they said to be “in balance” or “balanced”. Most of the audio inputs which purport to be balanced are like the two separate weights, they are unconnected inputs which, when a signal is applied to one, will produce an output which is 180 degrees out of phase relative to the other.

As we stated before, if the drive to each input is EXACTLY the same, and assuming the circuit has a big window before some kind of non-linearity upset the situation then you get zero output. But to achieve even 60dB of rejection you need calibration lab style resistors and components (and remember even the cable is part of the circuit) on both the transmitter and receiver ends of the transmission line. When you start to look at what a transformer can achieve, of the order of 120dB, which is 1000,000:1 then this scenario is just impossible. You could upset the symmetry of the circuit by just thinking about it.

Just when you thought the “actively balanced” circuit was leaning on the ropes with a bleeding nose and glazed eyes, it receives two more vicious blows.

Both relate to the amount of common mode signal that the circuit can handle before it starts doing something funky and unpredictable. The input circuitry of the op-amp or other circuit will have a (sometimes very limited) common mode range over which the component is linear. Outside that range and it becomes nonlinear and creates distortion. Nonlinearities at hum frequencies will upset the symmetry of the circuit and will drastically decrease the CMRR of the overall circuit. And of course the signal you’re actually interested in gets distorted as well. Normally we don’t expect to see a great deal of common mode hum voltage on the transmission line, if there’s a lot here, it probably means there’s a fault with grounding. It’s more of a problem at high frequencies, into the radio and microwave region, generated for example my mobile phones. At those frequencies the op-amp’s open loop gain will have dropped down to nothing, so the circuit doesn’t work properly anymore, That doesn’t mean that the circuit won’t respond to the RF, due to it’s nonlinearity it will demodulate the signal. There are various theories, which attribute “the transistor sound” to RF interference.

For example, a decent op-amp might have a slew rate of 10V/us that means that its output can change by 10v in 1 microsecond. That might seem pretty fast, but consider a 1GHz signal, mobile phone kind of frequency. At 1GHz, 1.6mV of signal is slewing at 10V/us. It doesn’t take a lot of it to completely overload and screw up such an active input which all manners of distortions resulting.

The transformer just doesn’t suffer this problem; a decently designed transformer component will have one or even two internal screens between the primary and secondary windings. These screens shield the secondary from the primary and prevent RF from getting across. The transformer’s natural high frequency limit (somewhat over 200kHz for most Audio Note™ units) prevents even differential mode hash getting to the input stage.

One final note, for optimum performance the source impedance should be low, and the load impedance, which the preamp for example sees should be high, especially to common mode signals. The 600 ohm source and load scenario is only relevant for true transmission lines used for example in telephone systems.

We used the single op-amp based circuit as an example, it is representative of the performance of any of the “actively balanced”, “symmetrical” and so on inputs, and there isn’t anything that can be done to improve that situation, not without making the whole thing far more complex, and unfortunately many/most manufacturers don’t even know how to do that, nor do they want to because it’s expensive. On top of that it’s fairly obvious to anyone who cares to listen that complexity is deleterious to sound quality. Those that for some reason think that complexity improves things should apply for a job at a large, well known software design establishment and abolish any thoughts .

The simulations give a graphic indication of what happens, it’s clearer than a pumping of numbers into a set of equations. Overall it can be seen that the transformer is truly and naturally a balanced input, and it gives superior performance to even the most “modern” and sophisticated op-amp based technology. All that’s happened over the years is that the art of designing a good transformer has died out, and that most modern electronic engineers just don’t understand them, so they use active electronic instead, often just cloning IC manufacturer datasheets, which is particularly tragic, and is just another example of the power of marketing hype in a world where few engineers actually do any real engineering, but prefer to rely on pre-digested information without checking its content with predictable results and go into complete denial when presented with these.

May 20, 2006
Andy Grove
Peter Qvortrup