

The G Word, or How to Get Your Audio off the Ground

Bruno Putzeys

My most daunting challenge is an ongoing one. I am trying to expunge my language of words that aggravate, cause hurt, misunderstandings or that are just meaningless blather. Fortunately I've never been one for sexual or racial epithets. Swear words were also easy to leave. Expressions like "dimwit" or "room-temperature IQ" are much harder to swear off. The one I find most troubling though is the G word. I can't bear to imagine saying it in polite company and yet all too I catch myself doing it unwittingly. Deep breath... I'm talking about... "GND". There. Forgive me. No more nastiness.

Audio signals are voltages. A voltage is the potential difference developed between two points. We grab a voltmeter and connect the two test leads to probe the two points, or "nodes" that we want to know the potential difference between. We don't just attach one lead and hope to get a reading.



Figure 1: Honestly, sir, I'm positive we had one of those in the school lab.

GND-think

And yet it is not unusual for audio engineers to think of an audio signal as only one circuit node or wire next to which a voltage is written or a waveform drawn, as though this single node were magically capable of having a voltage all on its own. The second node, it seems, is too unimportant, too obvious to mention. And this is where the rub lies: what on earth is *ground*?

According to GND Gurus the root cause of all hum and buzz problems is current flowing through "the same ground" as that used as voltage reference. So, they suggest, we use "different grounds".

The hidden assumption is that a signal is just one wire. But as anyone with a voltmeter knows, the second wire is every bit as important as the first. Still we seem to think it makes sense to use as the second wire the central sewage pipe that also carries waste electrons, supply return currents, shield currents etc back to the recycling plant. And then we're surprised to find rubbish on it.

The supposed solution is called a "star ground", a common point where "different grounds" connect.




Figure 2: It's got wonderful Powerpoint appeal, though.

It looks nice at first glance and its practitioners defend it as though it were a fundamental truth. Practically speaking though it's a nonstarter. It only works at all when it's rigorously done. You can star a power amp. You can star a preamp. And then you connect the two. Oops. Which of the two stars guards that mythical common potential that all signals in the combined circuit are referenced to? That's where GND Gurus get into their stride. Chains of stars, stars of stars, the whole celestial menagerie. All hinge on minimizing current flow through the connections that tie the local stars together. And so the saga continues with "floating grounds", disconnecting mains safety earth and whatnot.




Figure 3: Advanced GND Guruship in action. Yes, I found this on the Internet.

You heard me correctly. Most audio equipment has no safety earth connection simply because we can't seem to imagine signal connections without a common reference.

And often that doesn't even work. Suppose I have a TV, a DVD player and an amplifier. When I want to watch TV I want to hear the sound over my stereo. When I watch a DVD I'd rather run the audio straight from the player to the amp, not through the TV's rotten signal processor. So we connect the video output of a DVD player to the TV and the audio to the preamp and we also connect the TV's audio to the preamp. The dreaded "Ground Loop" scenario. Other than the most minimalist audiophile stereos there is no way of putting a system together without creating current loops. Current loops are a fact of life. Any scheme to avoid buzz and hum had better not rely on avoiding "ground loops".

The final nail in the star's coffin is that it *only works at DC*. A wire has inductance and two wires have mutual inductance on top of that. Accidentally lay a "dirty" return wire next to a "clean" reference wire and bam, noise. How do we add power supply decoupling? Do we run long wires from the decoupling capacitor to the star and add *exactly as much inductance as we were hoping to get rid of*? With a star you can just about build a mildly comatose class A amplifier. Anything faster and you'll run into stability problems. Try switching circuits and all assumptions go out the window.




Figure 4: How long PCB traces can a decoupling cap have before it's useless?

In short: any exposé that takes as its premise that hum, noise and distortion have something to do with "grounding" should be stamped on and ground into the ground. We need to design circuits that read voltages like voltmeters: with two wires. The result should not depend on the contents of the local electron tip¹.

Take-home message

When a change in "grounding" causes hum, this is because we're naïvely thinking of a signal as one wire. Stars are a band-aid to try to make this flawed assumption work.

Making a difference

Of course, such a way of working already exists. XLR connectors have an extra pin compared to an RCA connector. Pin 1 connects the chassis while pins 2 and 3 are the high and low terminals between which the signal is measured. Sadly enough, this too is riddled with confusing semantics, some of which have turned out to be deadly. Balanced, differential, symmetrical, what shall it be? Before I pitch my tent at any one of those three, let me quickly revisit what school books and audio magazines usually make of it.




Figure 5: The prototypical explanation of balanced connections.

The source, they say, produces two signals which are each other's mirror image. Any source of interference will affect both wires equally and the error is eliminated when the receiver subtracts the two signals. Note how the authors of this type of explanation have difficulty shedding GND-think. If those two signals are neatly symmetrical, about what potential exactly are they symmetrical? The source's return node? The chassis? Any of those on

¹ Note to non-UK readers, a "tip" is a rubbish dump.

the receiving side? And does it even matter? The input should only care about the difference between the two. The whole reason why the input measures the voltage between the two wires is precisely because *it's trying to ignore those irrelevant potentials*.

You can cut the amount of circuitry on the transmitting end by half simply by arbitrarily choosing some potential that it has handy anyway and connect one wire there. All it has to do is off-set the potential on the second wire to make the difference between the two the wanted output voltage.




Figure 6: Symmetry is useless.

This is just as good as the previous one. There's no pressing need to drive both wires actively. One will do. On the receiving end it's only the potential difference that matters. If one wire is connected to whatever node the source calls "my zero volts" the receiver duly subtracts the potentials of the two wires, regardless where its own personal zero volts might be with respect to the source's. I'm 1.8 m tall when I measure myself standing on the office floor. But this is equally true when I'm standing on a landfill. If you want to know my height, simply subtract the altitude of the refuse horizon from the altitude of my bald patch. There's no need for me to be dug in halfway.

This is seriously good news. To change an output from single-ended to differential all you need to add is an extra wire to carry the reference potential to the receiver. The burden then falls on the receiver to make the subtraction.

Take-home messages

Differential transmission of audio doesn't mean you need to make a symmetrical voltage.

An input that expects a symmetrical signal is not differential because it's trying to involve a third node into the equation whereas voltages are only measured between two nodes.

A word of terminology. The signal that we want to transmit is that measured between the two wires. This is also called the differential-mode signal. The error signal we want to ignore is the one that gets superimposed on both wires (as measured with respect to the receiver's chassis potential). This error signal may be due to interference en route, but in practice it's mainly the difference between the chassis potentials of the transmitter and the

receiver. That error signal is called the common mode signal.

Balance

The ideal differential input would be a transformer. By "ideal" I mean in terms of how well it would manage to look like a voltmeter with just two connections on it. Even if there were hundreds of volts between the chassis of the source and the receiver, this would go completely unnoticed.

Other than that, a balanced connection will look more like this:




Figure 7: A typical transformerless balanced connection.

V_{cm} symbolises any voltage between the two chassis, however it arose. If the input had been a transformer, no current would flow through the two signal wires, but transformerless inputs necessarily have some input network, if only to provide a path for base currents.

The task is to minimize the impact this current will have on the recovered audio signal.

Let's assume the source is putting out 0V and redraw the circuit as a Wheatstone bridge. Any signal seen between the inputs of the difference amplifier is unwanted.




Figure 8: Input/output resistances seen as a Wheatstone bridge.

$$\frac{V_{err}}{V_{cm}} = \frac{R_{ih}}{R_{oh} + R_{ih}} - \frac{R_{il}}{R_{ol} + R_{il}}$$

$$\frac{V_{dm}}{V_{out}} = \frac{R_i}{R_o + R_i}$$

It's clear that we don't need a transformer. We can allow current to flow through the signal wires so long as $R_{oh}/R_{ih} = R_{ol}/R_{il}$. If the input resistors are well-matched and so are the output resistors, no amount of common-mode voltage will get converted into an output signal.

When a Wheatstone bridge is exactly nulled, the term we use is that the bridge is *balanced*. That is where the word “balanced connection” comes from. It has nothing at all to do with one voltage going up while the other goes down, but with divider ratios being equal. Don’t think uppy-downy. Think equilibrium. Zen. Ooohmmmm...

The ratio between the error voltage and the common mode voltage is the common mode conversion ratio. The smaller it is, the better. It’s more common to quote this number in relation with the wanted signal, expressed in decibels. This ratio is called the Common Mode Rejection Ratio (CMRR)

$$CMRR = -20 \cdot \log \left(\frac{V_{err}/V_{cm}}{V_{dm}/V_{out}} \right)$$

Let’s explore for a second what happens if the output resistances are matched i.e., $R_{oh}=R_{ol}=R_o$ but the input resistances aren’t, say $R_{i1}=R_i$ and $R_{i2}=R_i+\Delta R_i$.

$$\left| \frac{V_{err}/V_{cm}}{V_{dm}/V_{out}} \right| \approx \frac{\Delta R_i \cdot R_o}{R_i \cdot (R_o + R_i)} \approx \frac{\Delta R_i \cdot R_o}{R_i^2}$$

The sensitivity to an imbalance in the input resistance increases with output resistance. It pays to minimize output resistance. It also decreases, quite rapidly, with increasing input resistance. So that seems a good idea too.

Secondly, let’s explore the impact of an imbalance in the output resistances

$$\left| \frac{V_{err}/V_{cm}}{V_{dm}/V_{out}} \right| \approx \frac{\Delta R_o}{R_i + R_o} \approx \frac{\Delta R_o}{R_i}$$

This is fairly important. If your input network consists of two resistors to some local reference, making those resistors as large as you can is going to make a lot of difference. And when you measure CMRR, do so with an imbalance of several ohms on the source side because that test will tell you a lot more about the real-world ability of an input to reject CMRR than a bench test with the inputs perfectly shorted together.

Acting locally

The biggest overlooked opportunity for differential signalling is *inside* the box where small signals and large currents slug it out in cramped quarters. Think about it. Class D amplifiers switch tens of amperes in a matter of nanoseconds mere centimetres away from where the line level signal comes in and gets processed and modulated. This is not an environment you want to try a star structure in.

Good EMI control requires a circuit board with one layer exclusively dedicated to serve as a solid copper fill or ground plane² which is devoted to the one circuit node which can acceptably be called GND. This is the one that supply and signal currents return to. All decoupling capacitors are directly connected to the ground plane, as are the “GND” pins of IC’s. The benefit of this relies on the fact that inside a conductor the current distributes itself inversely proportional to impedance.

When you trace all possible paths of a high frequency current through a trace and then back to the source through a copper fill, you’ll find that the impedance is mostly inductive, determined by area encircled by that path. One path is vastly more compact than all the rest: the one where the return current through the copper fill follows every turn and twist of the trace. If you force that current around a cut in the copper fill, this spot will become strongly inductive and develop a magnetic field. Cuts or splits in planes are an absolute no-no. Do not ever follow chip manufacturer’s layout guidelines if they recommend using separate analogue and digital planes or making cuts in the ground plane.

At low frequencies only resistance counts. Low-frequency currents will fan out widely over all available copper. Voltage differentials will develop all over the copper fill. A “ground plane” can’t be trusted to have the same potential everywhere. A correctly designed ground plane i.e., a contiguous one, is useless as a signal reference.

Do we worry? No. It just means we’re not going to use the ground plane as a signal reference. That should not be its function. Instead we’ll transmit every signal as a pair of wires. Now, would you believe you can do this throughout the circuit without adding active circuitry? Here’s how.

Step 1. The diff amp

In principle, you can make a differential amplifier using the classic difference amp circuit.




Figure 9: Badly drawn diff amp

² Not to be confused with that other “ground plane”, which is what engineers call the regular bus service run by Air France to ferry passengers between Brussels and Paris when French traffic control goes on strike again.

Wait. Ho. Stop.

There's something really wrong with this picture. Can you see it? Go ahead and see if you can spot it.

Here's what. What's the output signal in this drawing? Do we get out the magical uni-lead voltmeter again? We need to get serious about this. Every signal is two wires. Draw two.




Figure 10: Well drawn diff amp.

That's much better. We're getting into the swing.

You see, what this circuit does is amplify the input voltage by R_f/R_i and develop that voltage *between the output node of the op amp and whatever reference potential R_f is connected to.*

This is how it works:




Figure 11: Diff amp as reference translator.

This provides an alternative way of looking at the difference amplifier. It is a **reference translator**. It's a bit like a floating voltage source that you can reference anywhere you like.

But here's the shocker: you can add reference translation ability to any circuit, so long as you can build an inverting version of it.

Step 2. Generalized method

Suppose you have a circuit, e.g. a lowpass filter or, as the case may be, the loop filter of a class D amplifier. First, transform the circuit so that the non-inverting input of the op amp is tied to the reference potential. The block called feedback network can have several inputs. In this example it has two: one is the signal input, the other the feedback input. It could equally have multiple signal and feedback inputs. Finally it may have a connection to the reference potential.




Figure 12: Generalized inverting circuit.

Once you've got that, flip the feedback network over.




Figure 13: Differential execution of inverting circuit.

That's all there is to it!

The node that was originally tied to the ambiguous node called "ground" no longer needs to be connected anywhere unless required to keep the op amp input from overloading with common-mode signals. In that case the sensible place is on the ground plane near the op amp's decoupling caps, as this is implicitly the HF reference of the op amp. Ideally the feedthrough from this point to the differential output voltage is zero, which is why we have that liberty.

We now have a proper differential pair. One wire is actively driven by the op amp, the second one is passively driven by a low-impedance tie to the ground plane that can be made pretty much anywhere. All that matters is that the whole trace has just one such connection so that the next stage, which you have also transformed in this manner, takes its input between the same pair of nodes as the feedback network. Always route signals as two wires, one right next to the other to minimize magnetic pick-up and balance capacitive pick-up. You might need to make the passive drive connection through a zero ohm resistor to make sure your PCB layout software understands that the second wire is to be treated as a separate net, even if it is galvanically connected, at one point, to the one called GND.

Step 3. Let the Strong Help the Weak

Some circuits don't have a virtual-ground pendant. Circuits with potentiometers, in general, do not lend themselves well to this approach.

In that case, we can use the level-shifting capability of the surrounding circuitry to solve the problem.

Imagine a problem circuit flanked by two differential circuits.




Figure 14: Extant stages as problem solvers.

The problem circuit has a single reference node which it uses for input and output. Tie this node to the ground plane at a single point and use that for all "GND" connections of this particular subcircuit. The output of the first stage and the input of the second stage are both made to reference this point.

I think we can see how this solves the grounding question rather magnificently. When you have a chain of signal processing stages laid out on a board with a solid ground plane, each signal run between stages is referenced at the most convenient point on the ground plane. There is no reason to try and make a global reference. The differential signal just hops from one reference to the other as it progresses through the circuit.

Impedance Balance vs. Current Balance.

A confounding aspect of diff amps is that the input currents are almost never equal. Here's the situation. If I hold one input of a 10x diff amp at ground potential and drive the other, I get a factor 11 discrepancy in the input current, depending on which is the one that gets driven. And so, the unsuspecting engineer might naively reason, the input impedance is out of balance and should be put right.




Figure 15: The Imbalance Illusion

What they do in response to this misconception is quite ghastly.




Figure 16: Grisly outcome of cognitive illusion.

You can easily see why there is something fishy about this. If instead of driving the circuit with one leg grounded we drove it with a symmetrical signal, the ratio between the input currents would no longer work out as 11:1 but as 21:1. You can't scale the impedances in a way that the currents work out equal under all conditions.

What's going on here? Remember to think of a differential input as forming a Wheatstone bridge along with the source resistances. If you add source resistors to the above circuit you get something that is clearly no longer a difference amplifier.




Figure 17: Why it's grisly.

We should repair the circuit and make the 2 legs equal again. We once again have a fully functional diff amp. If the input currents are different, this is no indication of imbalance.




Figure 18: Wei Wu Wei, or how balance is restored by not intervening.

We should have seen from the start that the problem was illusory. In order to contrive it we had to drag the output reference of the difference amp into the equation *and falsely assume that this is the point that the common-mode input impedance*

refers to. The circuit is balanced alright but it just so happens that the input impedance references the virtual short, not some handy point that someone calls GND.

Take-home messages

Converting a circuit to differential does not require additional amplification stages.

Each signal has its own reference.

Making a circuit differential is not the same as building two independent copies of a ground-referenced one.

Do not try to equalise signal currents. It doesn't work and you'll end up creating an impedance imbalance of heroic proportions.

In From the Cold

As a means of immunizing your circuit to circulating currents in the ground plane, difference amplifiers and more generally differential circuits are a smashing idea. As a means to building a robust interface with the outside world they need an extra ingredient.

We mentioned earlier that the sensitivity to common-mode errors depends strongly on the input impedance. The lower it is, the more crucial matching will become. After all, a low input resistance will convert any common-mode voltage into a common-mode current, and any matching error will then go on to convert that into a differential-mode voltage at the input.

Diff amps, for reasons of noise, are low-impedance affairs and hence their balance is easily upset by a source impedance imbalance. For connections to the outside world it is good practice to buffer the input signal.




Figure 19: The instrumentation amplifier.

The circuit we get is called the instrumentation amplifier. This circuit is usually drawn with the first stage not just buffering the input, *but providing all of the gain as well*. There is a significant incentive to doing so. Note how, regardless of any mismatches of R_{f1} and R_{f2} , the first stage will never

convert common mode into differential mode. The common mode component is passed through unchanged but only the differential mode component is amplified. The ability of the second stage to tell the two apart is multiplied by the gain of the first stage. Gain in the buffer stage adds free common-mode rejection.

I can see certain sections of the readership bristle at this point. Isn't adding an extra amplification stage in the signal path worse than the illness? Well, if you're of that penchant I can only say: try it. You'll discover that modern op amps change the sound a lot less than the noise added by a single unbalanced connection. There is a reason for the sprawling cottage industry that audiophile cables have spawned. Unbalanced connections affect the sound quite strongly for reasons that by now should be quite obvious. The whole idea of having one of the actual signal wires also do the dirty work of shunting equalizing currents away is barmy. To then try and solve the problem by eliminating those currents is bone-headed. To try and mollycoddle the sonic defects this causes by making outlandish cables is madness. The RCA connector and all it stands for should be banned by law.

Where was I? Oh yes, instrumentation amps. It would be a wonderful world where we could just connect the noninverting inputs of the first pair of amps to the outside world. This is not quite practical because they tend to leak. This is why input resistors were assumed. The ones we wanted to make as large as possible to keep common-mode voltages from becoming input currents. Let's see. There could be two ways of looking at them. Ideally we'd want them to be low enough so that when the input gets unplugged the bias current of the op amp inputs doesn't produce a DC shift big enough to get a thump out of the speakers. Looking at the data sheets of commonly used op amps that would still only be in the 10k's. Or we could lessen the requirement and ask merely that the bias current doesn't cause the input to drift off more than a few volts. After all, the source might be AC coupled.

Well, we can get both right. The signal that thumps the speaker is a differential DC component. Strapping a sufficiently small resistor (tens of kilo-ohms) across the two input terminals will fix that. We can now afford to allow both inputs to drift off-centre by several volts if need be. It's common mode, so what. So let this resistance be large (megohms). It's effectively the common mode impedance we need maximized, not the differential mode one.




Figure 20: Optimal input biasing.

What we've done really is to insert a large resistor R_{cd} in series with the receiver side ground leg of the Wheatstone bridge of Figure 8. This resistor will greatly reduce the current flowing through the signal wires. Calculating the exact impact on CMRR is left as an exercise to the reader. Meanwhile, the two input resistors insure that the bias currents of the two op amps do not result in a too large differential mode DC voltage, even when one terminal is tied down and the other left floating.

The same arguably holds for the source end. If the output can be floated, this too will limit conversion of common mode voltages into current. But while it's trivial to achieve common mode input impedances in the meg range, other than using a transformer, doing the same on the source side is very daunting indeed. Unless there are pressing reasons to build an electronically floated output I shouldn't bother.

The elephant in the room here are input filters. While virtually unknown in high end consumer audio, input filters are invariably added to properly designed kit to insure that the music continues in the presence of mobile phones and taxi dispatches. Input filters are expected to block RF as it tries to enter the enclosure. They comprise capacitances of 100 pF or more, directly connected to the chassis. They are an integral part of the Wheatstone bridge, are rarely matched and drive down the common mode input impedance. 2×100 pF at 20 kHz works out as about 40 k. That's a far cry from the megs we can get at DC. The problem has been quite elegantly solved by Bill Whitlock³, who uses a bootstrapping technique to increase the common mode impedance of the input filter. Do look it up. Unfortunately the only way to use this method is to buy IC's from the current licensee of his patent.

³ Whitlock, Bill, A New Balanced Audio Input Circuit for Maximum Common-Mode Rejection in Real-World Environments, AES pre-print 4372

Take-home messages

The unbalancing effects of mismatch in the source resistance are exacerbated by low common-mode input impedances.

A noninverting differential gain stage allows very high CM input impedances and reduces matching requirements in the diff amp stage.

Wiring Up

Balanced audio cables are shielded twisted pairs. You can often make a perfectly good connection with an unshielded twisted pair provided all boxes live roughly at the same potential (to keep from overdriving the common-mode voltage range of some inputs) but "often" is not good enough in the real world.

So far I've treated differential connections like they should: as two wires. Confusingly, XLR connections have three pins. Pins 2 and 3 are the noninverting ("hot") and inverting ("cold") wires. So far so good. Pin 1 though, is designated "ground". By now we know to ask the question: what on earth is meant by such an ambiguous word?

In spite of much confusion there is a clear and unambiguous answer. It should be connected in a way which allows the shield (the braid that goes around the signal pair) to perform its function. That function is to make a tunnel-like extension between two chassis inside which the actual signal pair is well-protected. That's what "shield" means.




Figure 21: What the cable shield is for.

Ideally we want the shield to bond directly to the rim of a circular hole in the chassis on both ends. If this is impossible, try to get as close to this ideal as you can. To get an idea of the effectiveness of a cable shield, consider this: the current through a hollow conductor does not create a magnetic field inside that conductor. All of it is outside the shield and hence around the signal wires as well.




Figure 22: CM reduction effected by shield

It's the same neat trick that makes coax cables work. A high-frequency current through the shield induces the same voltage along the inner conductor as along the shield. At high frequencies the input voltage of a coax cable (as measured between the conductor and the shield) is the same as the output voltage. By the same token, the shield serves to reduce the common-mode voltage at the receiving end of a balanced audio cable.

How high is high? Well, practically speaking, for a normal braided shield and a cable of a few metres long this effect starts becoming noticeable from a few hundred Hz upward. Below that the shield is still a pure resistance and the full voltage differential along the shield appears as common mode.

For this to work the shield should be bonded to the chassis with the lowest possible impedance. The shield itself should be a neat cylinder round the signal pair and have low resistance too. A foil shield with a drain wire is a no-no, because the drain wire concentrates the current onto itself and defeats the common-mode reducing effect up to frequencies well above the audio band. In fact, it may even preferentially couple noise into one of the two signal wires, thereby turning shield current into a differential mode error. This effect is known as Shield Current Induced Noise (SCIN)⁴.

The Pin 1 Problem

The XLR connector has been something of a missed chance. It should have been a round shell with just 2 pins in it. Nobody would have doubted that the shell should connect at the chassis. But now it's got the third pin which has misled people into thinking that it was some kind of "audio ground" connection that should connect somewhere other than the shell.

What happened is that a lot of people connected pin 1 to their internal zero volt reference (infelicitously called GND). Instead of shunting away circulating currents into the chassis, this actually invites them in to have an all night party, romp around in the furniture and be sick all over the carpet.




Figure 23: The Pin 1 Problem: follow the white - the current.

Circuits designed according to the differential method explained earlier are insensitive to pin 1 problems. The PCB copper fill functions more like a chassis than a reference. Later in our demo project we'll be cheerfully tying pin 1 and the shell to the ground with no ill effect. But this is not how most equipment is designed. Most are single-ended internally and they do use ground as a global reference. "Ground current" means current through your internal reference. Instead of being a handy point to bond the cable shield to chassis, pin 1 has inadvertently been turned into an input.

Pin 1 problems drive users mad. The trouble is that they affect input and output connectors equally. You can induce hum in many products by feeding current into pin 1 of an output. So you could be fooled into thinking that "there's no hum on the output" because there's only hum when it's connected to a specific input. And that input gets the blame.

Oftentimes the blame is placed on "ground loops". Well DUH! That's like blaming a broken teapot on gravity. Cases where you have no loops are so rare that they're almost accidents. Circulating currents in audio cables are mostly unavoidable. It's just plain good engineering practice to make equipment immune to them. At some point the problem became so prevalent that the AES had to enshrine the obvious into a standard. Called AES48, it patiently explains that the shield should be connected to the chassis via the shortest possible route and that connections between the PCB ground and the chassis should be made elsewhere.




Figure 24: The right way to connect pin 1

Fortunately that's all you need to do to solve the problem. Informed studio techs don't even bother hunting down hum. They simply open up every box right after delivery and modify it so that it is AES48 compliant. Exit hum.

You should take the exhortation to keep the connection between pin 1 and chassis short very seriously. Cables get subjected to GSM radiation and worse. If the connection between pin 1 and the chassis is a long piece of wire, that radiation is now inside your chassis.

⁴ Brown, Jim; Whitlock, Bill, Common-Mode to Differential-Mode Conversion in Shielded Twisted-pair Cables (Shield-Current-Induced Noise), AES preprint 5747

Take-home messages

The cable shield is an effective tool to reduce common-mode noise.

In order to work it has to be a nice, round cylinder with the signal pair neatly in the centre, and it needs to be bonded straight to the chassis on both ends.

We've learned a lot. Let's put a few things in practice.

Demo project: a balanced volume controller.

According to a quick scan of professional audio forums, a perennial question is how to build a purist balanced volume controller. Two recurring themes are H pad attenuators and dual gang pots. H-pads attenuate the differential mode component without affecting the common mode component. At low volume settings, effective CMRR of the whole system may even become negative. H pads are out. A 2-gang potentiometer will convert CM to DM unless matching is phenomenal. Other than that, CM impedance is directly determined to DM impedance. For noise and distortion reasons you'd like a low resistance pot, for CMRR reasons you'd like high resistance. This is going nowhere either. It turns out that there is no acceptable method of constructing a balanced passive volume controller. In fact, there is no sensible way to arrange a potentiometer in a differential fashion. This is where we use the tactic shown in figure 14. Since the volume controller stage has only one reference common to the input and the output, the surrounding stages are recruited to do the translation.

I have a double agenda in presenting this demonstration project. Firstly just to demonstrate how the "new" design methodology works in practice, but secondly to invite doubters to discover for themselves how a bit of rational engineering can produce staggeringly good sonics without resorting to boutique parts or boutique thinking. This is going to be the cheapest and best-sounding preamplifier you've ever built.

The input stage

The input stage is a straight buffer implementing the improved input biasing network. I would've used the Whitlock's input chips and implemented the capacitive bootstrap technique as well, except that the distortion performance is not good enough in my view.

The difference stage

As said, we're out of luck when it comes to wiring a pot differentially so we won't even try. Instead we'll be using the surrounding stages to reference the

cold point of the variable gain stage. So between the input buffer and the variable gain stage we insert a difference amplifier. This is the circuit that'll confer CMRR to our little preamp, so resistor matching is of prime importance here. The output of the difference amplifier is referenced to the cold point of the volume controller.

The DC servo

I've always considered it the task of the preamplifier to remove DC. I've thrown in an unusual DC removal circuit that isn't actually a servo in that it doesn't measure DC at the output. Instead it's a 2nd order low-pass filter whose output is subsequently subtracted from the signal.

The volume controller

As most experimenters will have noticed, potentiometers leave wildly varying and occasionally unpredictable footprints on the sound. Postmodern etiquette then requires that one congratulates oneself on having heard something that doubtlessly the objectivist clique will never have noticed and will most certainly deny, so the new observation is set aside for wonderment and mysticism and, crucially, exploitation by manufacturers of very expensive parts.

I must disappoint the postmodern set here. The problem is perfectly well known and well-understood if not by too many people. There are two elements at play. The resistive track is rarely linear. On top of that the non-linearity is dependent on the current density in the track. In logarithmic pots the divider ratio becomes non-linear. Also the wiper contact is a source of distortion. Secondly, very few amplifier circuits have a perfectly linear input impedance. It doesn't even matter whether it's valves, JFETs or bipolar, op amp or otherwise. All have, to a lesser or greater extent, a variable input capacitance. Drive an amplifier circuit with a few kilo-ohms at your peril.

Two exceptions. Virtual-ground circuits have no input capacitance modulation problems because the input signal is zero. Differential circuits have no problem either because the nonlinear charge currents cancel.

Whoa. Not only does differential circuit design do away with current loop problems, it actually eliminates a significant source of distortion. If panaceas exist, this must be one of them.

Long story short. Instead of operating as an attenuator the potentiometer is used as the sole feedback element in an inverting amplifier. Linearity of the volume control now only hinges on the linearity of the divider ratio. This is almost guaranteed in linear pots. The track resistance can be very very non-linear before this becomes an issue. Just

to make a point I decided to use a cheap 9mm “car stereo” pot. Distortion performance is top notch. The only drawback is that the control law follows an S-curve. At both extremes control becomes very sensitive. Since preamps are rarely used beyond unity gain you’ll mostly find the it to be a bit fiddly at quiet settings. We’ll have to live with that because adding external resistors to modify the control law will immediately put the linearity of the track resistance back into the equation. As it is, channel matching is surprisingly good even down to moderately low settings.

The output stage?

There’s no output stage! Well, there is, in a way. They’re the two 22 ohm build-out resistors whose only function is to isolate the cable capacitance from the op amp. Referring back to Figure 6 there’s no point in doing anything with the signal other than to provide connections to both ends of the signal i.e., the output pin of the op amp and the potential that the variable gain stage calls “my zero volt reference”.

PCB lay-out.

Differential circuit design treats every signal as a pair of wires. Usually though, only one of them is actively driven. The other is tied to the ground plane at some point by means of which the two processing stages at either end agree to call this particular potential “reference potential”. Circuit board layout programs tend not to like this kind of thing. When a connection is nominally the same net as the ground plane, they’ll nail every pin to the ground plane at the slightest excuse. And if you do make it a separate net, anything you do to make a galvanic connection to the ground plane is treated as a design rule error. With some layout tools there is nothing for it but to use a physical zero ohm resistor to connect the nets together when the board is assembled. Others, like Altium allow a kind of “part” called a “net tie”. When a part is declared to be a net tie, short circuit checking is locally turned off for that part allowing you to make overlapping pads. That’s what I’ve done here. My net ties are clearly recognizable in the board layout as two overlapping circular pads to serve as a visual aid to see what’s going on.

A salient feature of the board layout is that all components are placed in pairs. This runs counter to the usual practice of making the hot and cold sides of a differential circuit each other’s mirror image. But remember what we’re trying to do: we’re trying to make sure that any interference affects both legs equally. Another way of putting this is that the area enclosed by a differential pair must be as small as possible. Mirror-image layouts are exactly the wrong way to do it. If you want to have

a visually appealing symmetry, do so with the left and right channels.

Power supply

As a wink and a nod, the power supply is based on my HPR12/HNR12 regulators. The foot-print is compatible with ordinary 7812/7912 parts though.

Test results

I tested the circuit with a 600 ohm load and started with a 1kHz THD+N level sweep. Both at unity gain and -20dB, onset of clipping is just above 19dBu (6.9Vrms). Clearly the difference amp stage clips first.




Figure 25: Distortion+Noise as a function of input level at 1kHz.

Other than indicating the maximum signal level this plot doesn’t say much. Noise is clearly visible at lower levels but at higher levels the reading is dangerously close to the noise floor of the analyser. A THD vs frequency sweep was more revealing. For this test I set the analyser to measure just the harmonics and ignore the noise. The input level was set to 18dBu which is pretty close to the clipping point and quite a common choice in professional equipment.




Figure 26: THD without noise as a function of frequency at 18dBu

The rise at low frequencies is attributable to uneven thermal modulation of the resistance track. In retrospect I should have picked something like a Cermet pot. In spite of this, you will find it very difficult to find a cleaner preamp, regardless of price or fancy parts. Note the absence of distortion at the top end of the audio band. Any form of capacitance variation would have resulted in a rise with frequency.

There was no measurable difference between 100k loading and 600 ohm loading, even though the latter forces the op amp far into class B. This demonstrates the complete indifference of the circuit to distorted currents returning through the supply lines and the ground plane.


Sensible listening tests

When you decide to build this preamp, build two. That way you can use the second preamp as a volume controlled A/B switch to compare, variously, an expensive high end preamp (set to unity gain),


this little preamp (also at unity gain) and a direct connection to the source. Listen to which of the two preamps' output resembles the input signal most. You may find the experience enlightening.

Drawings


Schematic




PCB, component placement



PCB, top copper



PCB, bottom copper



Parts list

Designator	Part description
C1, C2, C3, C5, C18, C19, C20, C22	100p 50V NP0 0805
C4, C7, C8, C9, C21, C23, C24, C25	10u 50V Non-Polar
C10, C16, C26, C27, C28, C29	22u 63V
C12, C13, C14, C15	1000u 25V
D1, D2, D3, D4	RS1D
D5	BZX384C18
D6, D8	BAS316
J1, J3, J7, J9	NC3FAH2
J2, J8	NC3MAH1
J6	JST-B3P-VH
K1, K2, K3	Relay 12V 2x1A DPDT
R1, R9, R11, R18, R29, R36, R38, R45	100R 0805 thin film
R2, R3, R13, R14, R30, R31, R40, R41	1k 0805 precision thin film
R4, R15, R32, R42	22R 0805
R5	PTD902-2015F-A103
R6, R12, R33, R39	47k 0805
R7, R16, R34, R43	12k 0805
R8, R17, R35, R44	10k 0805 precision thin film
R10, R37	2M2 0805
R19, R22, R46, R49	10k 0805
R20, R23, R47, R50	220k 0805
R24, R54	2k2 0805
R25, R26, R27, R28	2R2 1206 high current
S1	Toggle switch SPDT
T3	BC846B
U1, U2, U6, U7	LM4562 SO8
U3	TL072 SO8
U4	7812 or HPR12
U5	7912 or HNR12