http://peufeu.free.fr/audio/memory/

Memory Distortion

This page is about a little known form of distortion called Memory Distortion, which happens in solid state devices. It is a bold attempt made to explain the "solid state sound", or the "sound of feedback".

Memory Distortion was discovered a long time ago by a French man called Gérard Perrot, who has since created his own company called Lavardin, which makes amplifiers and other Audio devices. I have tried to explore his ideas by building a different kind of amplifier. Here are some of my experiences.

Collected articles from other people :

Very interesting article that shows how to actually *measure* these elusive thermal drifts (hint : it is based on the integrated dc error over time). [Download]

Patent on the global amplifier design. Recommended reading. [Download]

Patents on the Lavardin input stage : interesting, bur read the other ones before that one.[Download]

Memory Distortion - Part 1 : Theory

Copyright Advice : I first heard of this in a patent from a French designer named Gérard Perrot so Thanks to him. He owns intellectual property on some of the circuits. Fortunately this only concerns commercial matters and DIYers can't get sued from building stuff at home, even if it is patented by other people...

0. Introduction

- **Fact A:**

A French audio manufacturer, **Lavardin**, claimed the discovery of a previously neglected circuit characteristic, which they called "circuit memory". The root of phenomenon was, according to them, signal-induced thermal drifts in transistors. They developed an amplifier which supposedly eliminated this phenomenon.

- **Fact B:**

Said amplifier got reviews to die for. Audio critics got hysterical. JM Lab uses it to power their Utopia line of loudspeaker in demos. However, the manufacturer in question being French, having limited production capacity, and expensive prices (although no so by audio standards), the scandal somehow faded as few people (not including me) could test listen the amplifier. Lavardin, however, seems to be doing good business.

- **Fact C:**

When there's smoke there's fire.

What Now ?

It is not often that some claimed "discovery" in audio makes sense. I got a copy of the patents. They made sense. I'll try to sum up what I understood. I don't pretend to be the owner of the absolute truth. I do think, however, that I might attract your attention on a previously "neglected circuit characteristic", tickle your curiosity, and maybe interest you.

1. Definitions

First, the obvious : circuit memory is the ability of a circuit to remember the past states of a signal. Here are some examples of memories :

- A capacitor has memory : it integrates current into voltage
- An inductor : it integrates voltage into current
- A heatsink, or a transistor : its mass stays warm long after thermal power has been dissipated.
- A low-pass filter
- A transistor driven into saturation
- etc.

I will not talk about inductors and capacitors directly, as integrating is part of their function. That's how filters are made.

Stored charges in junctions : this is a known problem, of which I am no specialist. They determine how long the amplifier takes to get out of clipping ; this is important because as long as the amplifier is in clipping, all signal is cut off. This is not of this kind of memory that we are going to talk about, however.

2. A striking example of Memory, or Why Distortion Measurements Fail

Check out this simple, bare circuit :

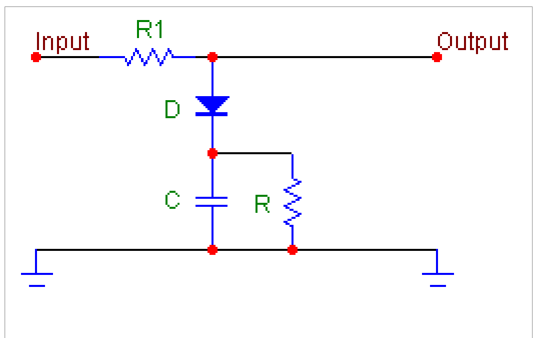

Fig.1 : the THD Trick

Now imagine we measure its distortion, by sending a 1 KHz sine of constant amplitude at the input and connecting the output to an analyzer. With adequate component values, after a few seconds the cap is charged with the peak voltage of the sine wave minus a diode voltage. Once it is charged, almost no current flows in the diode anymore, and the sine wave passes untouched, giving very good distortion measurements.

If we were to listen to music through this, however, we would notice that, after several seconds of low amplitude signals, the cap has discharged and will completely clip the crescendo that follows. This circuit has memory : the cap voltage is an image of past amplitude peaks.

This example is rather extreme, but shows that measurements using constant amplitude signals will fail to describe a system that is used to reproduce signals with extremely wide dynamics !

3. Distortion and Operating Point

Check out this other, simple circuit : a one transistor amplifier. It is a simplification of the Voltage Amplification Stage (VAS) of a solid-state amplifier without current mirror in the first stage. We are going to examine the effect if operating point variations.

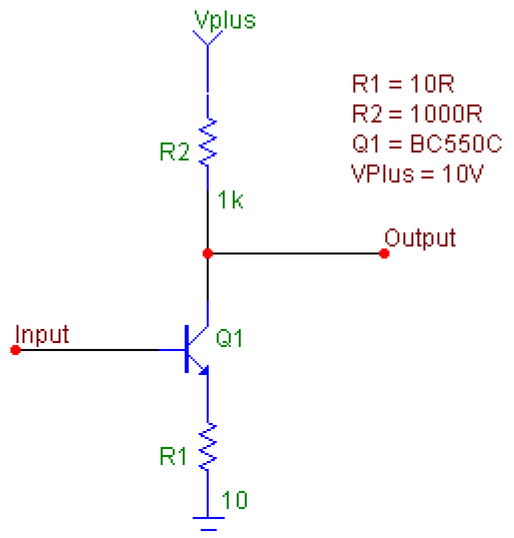

Fig. 1-2 : Simplistic amplifier

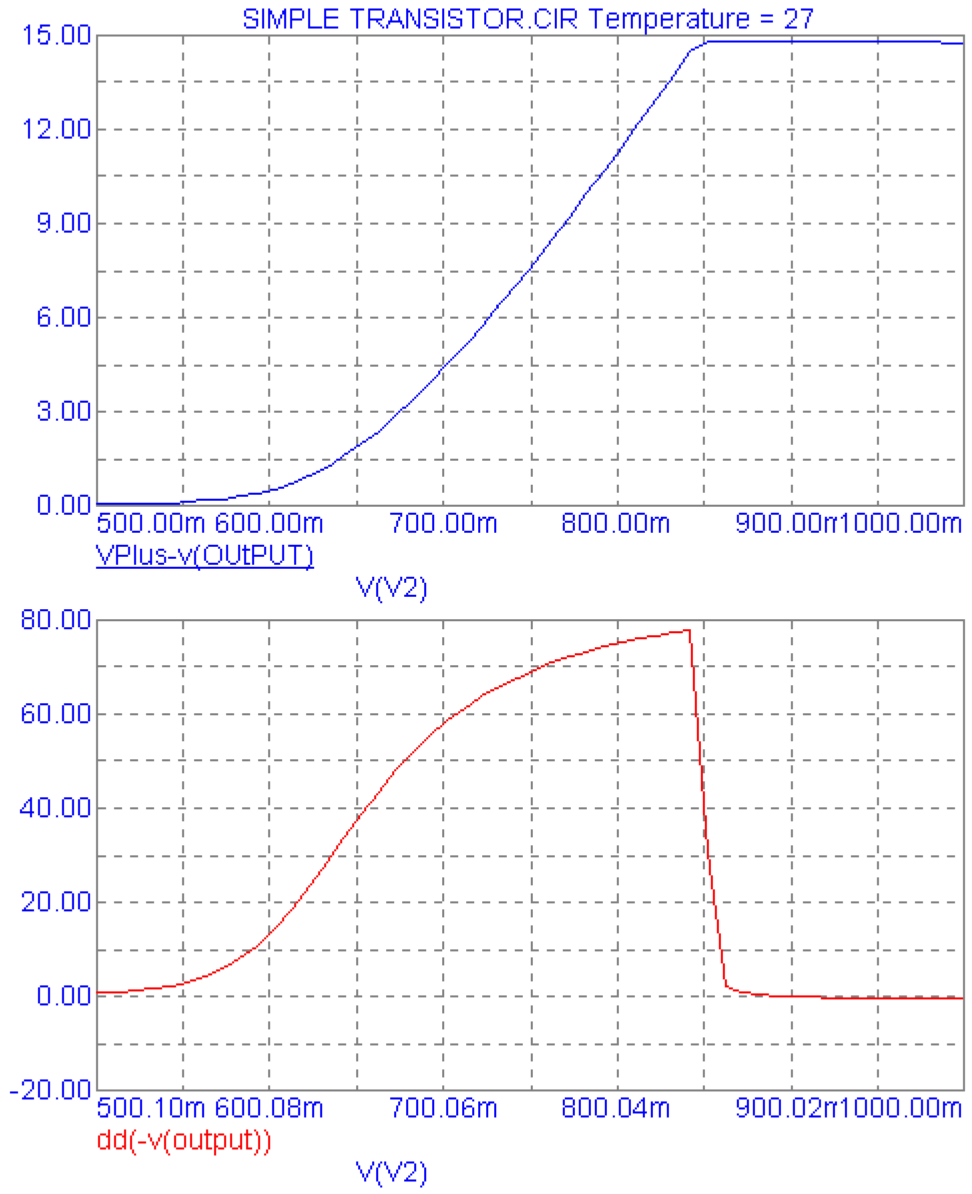

Fig. 1-3 : Output voltage (inverted, upper curve)
and derivative (lower curve) versus Input voltage

Fig.2 shows the transfer function of this amplifier : a portion of exponential fading into a straight line as the emmiter degeneration resistor kicks in. This particular circuit is not very linear, as the gain curve shows (red curve in Fig. 1-3). We see clipping, also.

If we were to use this amplifier in real life, we would need to choose a DC bias point around which to operate. Let it be VInput=0.7V+AC. Now, we can re-center the transfer function curves around the DC bias point : Fig. 1-4 has the transfer function for an input signal of +/- 20 mV around the DC Bias.

I also stepped the DC Bias from 700mV to 750mV in 5mV steps : these are the series of curves on Fig. 1-4. Transfer function is still the blue one and gain (derivative) the red one.

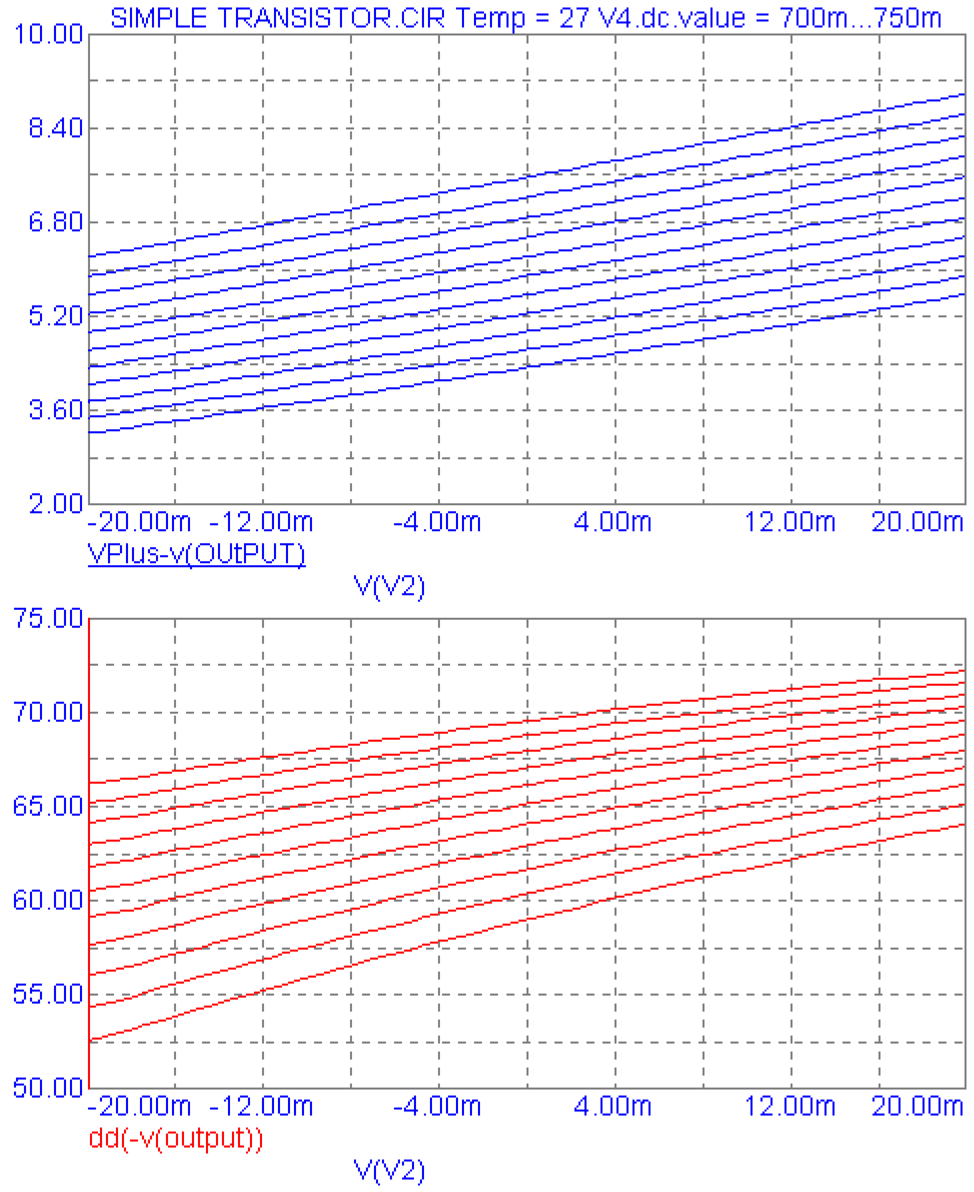

Fig. 1-4 : Effect of DC Bias variation on transfer function

Shifting the DC bias simply shifts the curves horizontally, their *global* shape does not change at all. However, the part of the curve that is used for amplification is the one at the **center**, therefore its shape *changes* with DC Bias.

Each of these transfer function has its own non-linearities, gains, and probably sonic character. It is safe to assume that if we built the circuit and measured its distortion, the harmonics spectra would be different between 700mV and 750mV biasses, and even between each 5 mV bias curve.

Besides, if we use this circuit in an amplifier with global feedback, which is typically done, its gain will be a part of the global open-loop gain. The open-loop gain will then be dependent on the DC bias of that transistor.

Executive summary : Every parameter of the transistor depends of its DC operating point (Bias) :

- Gain
- Non-linearities (distortion spectrum)
- Possibly, Sonic character

So, am I re-inventing intermodulation distortion, you will ask. For all intents and purposes all I wrote above in 3. is strictly a rehash of known stuff on intermodulation distortion (IMD). Besides, IMD is not a problem in a well designed, linear amplifier. So what ? Wait until you read the rest...

4. But DC conditions are fixed, aren't they ?

Consider the following simplified amplifier embryo (Fig. 1-5) : a diff pair (Q1 Q2) and a VAS (Q3). I put a PNP diff pair, a la Douglas Self, for a change.

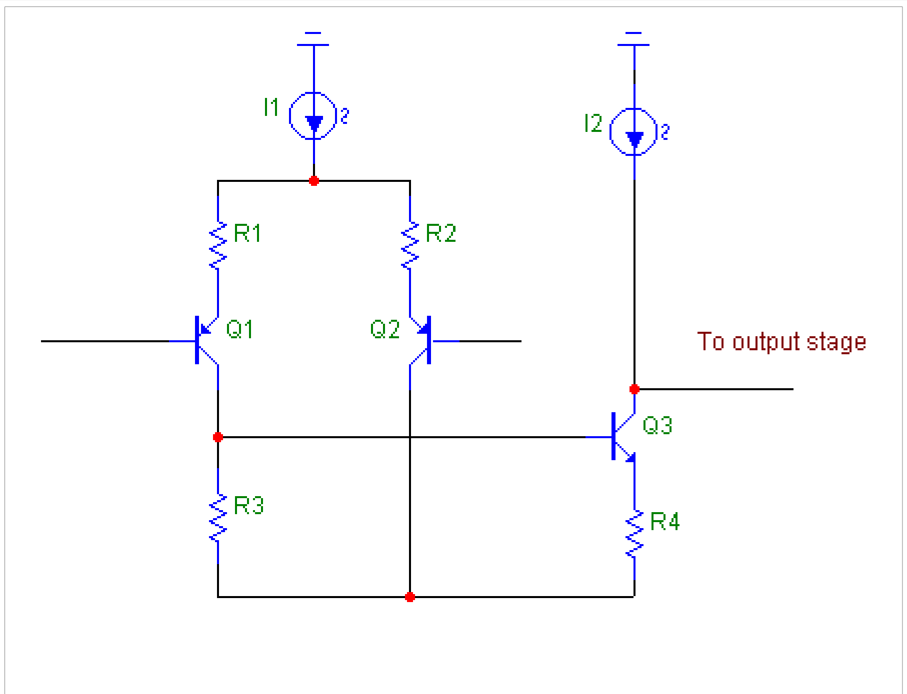

Fig. 1-5 : Some bits from an amplifier

This makes for a high feedback design, with R1=R2=10R, R3=1K, R4=10R. Compensation cap is not drawn.

Let us do a thermal modelling of a realistic VAS transistor (Q3) : it will be a TO126 critter (BD139 as the usual suspect). Voltage-driven (like above), with I2=20mA and R4=10R. Typical values. The collector is loaded by a buffer transistor. Our thermal model is in Fig. 1-6 :

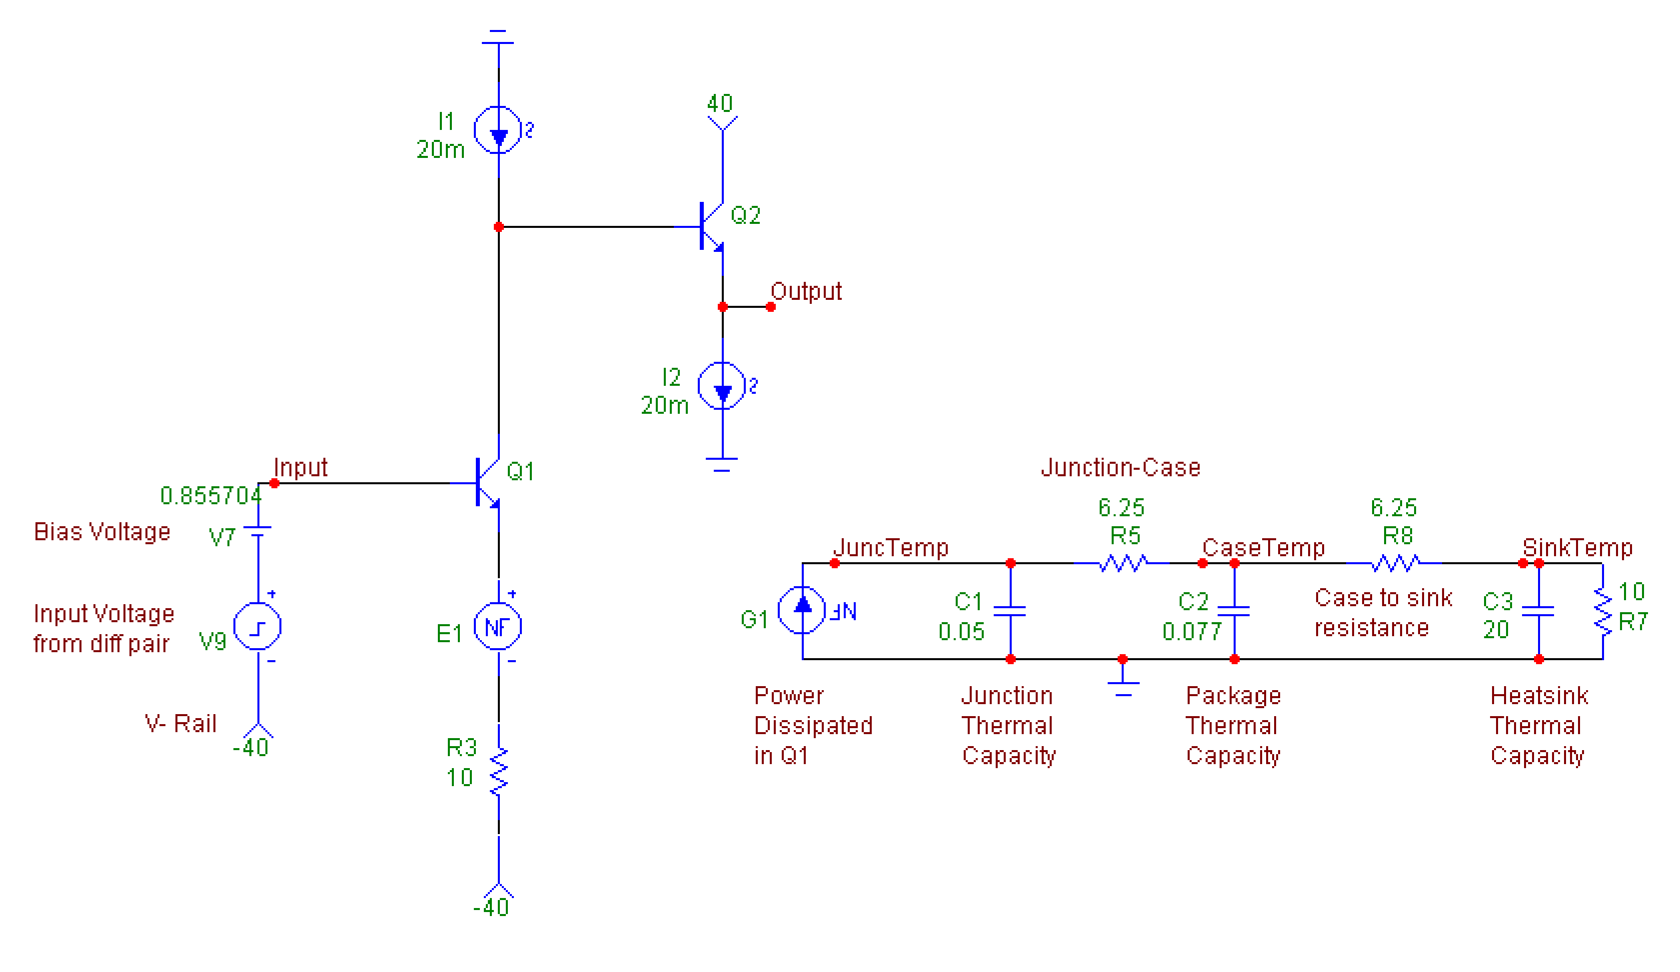

Fig. 1-6 : Thermal model

Notice G1, a current source whose value is proportional to the power dissipated in Q1 (Vce*Ie), and the various caps and resistor that model the thermal capacitances and resistances of the transistor package and heatsink. Notice also E1, which models the Vbe variations of -2 mV/Celsius.

Now, let's drive it with a large transient signal, as shown in Fig. 1-7, which does not drive it into clipping, but close, and plot the temperature variations of the transistor silicon. This transient signal is 30V at the output for 20ms. The output voltage is in blue, the junction temperature in red :

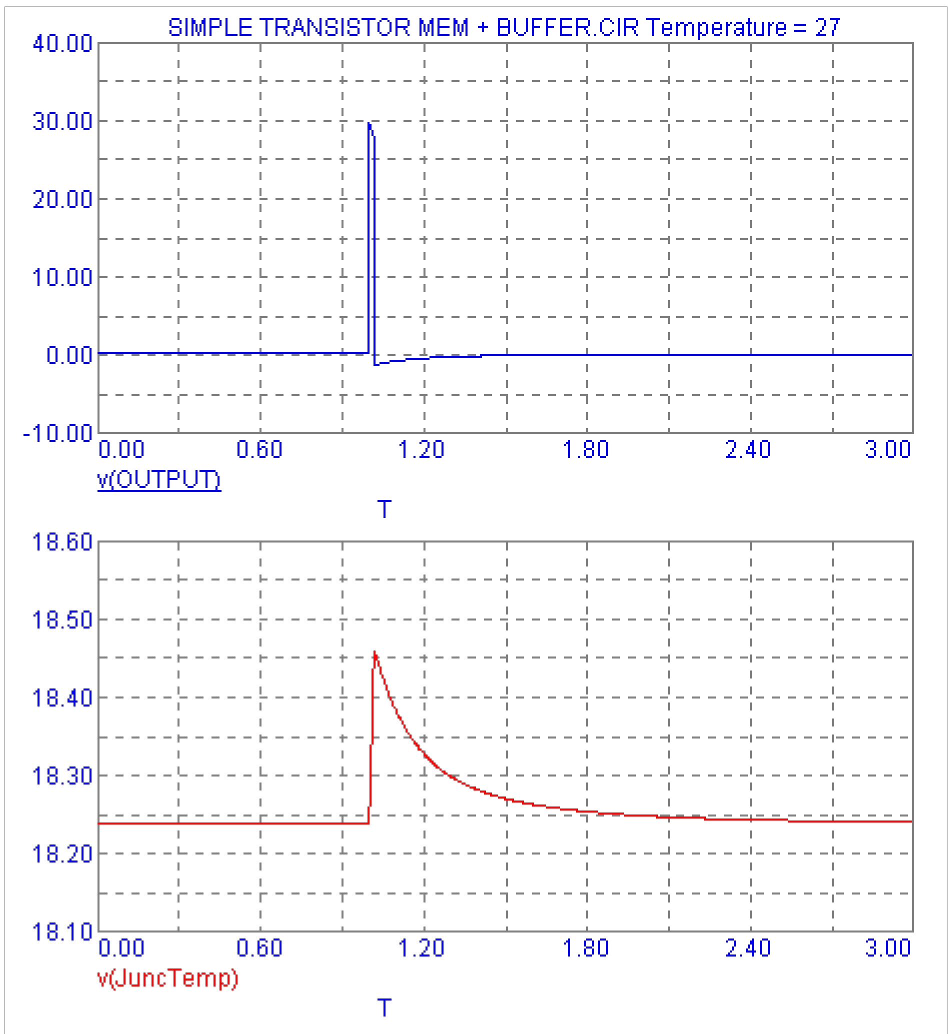

Fig. 1-7a : Temperature variations on a large transient

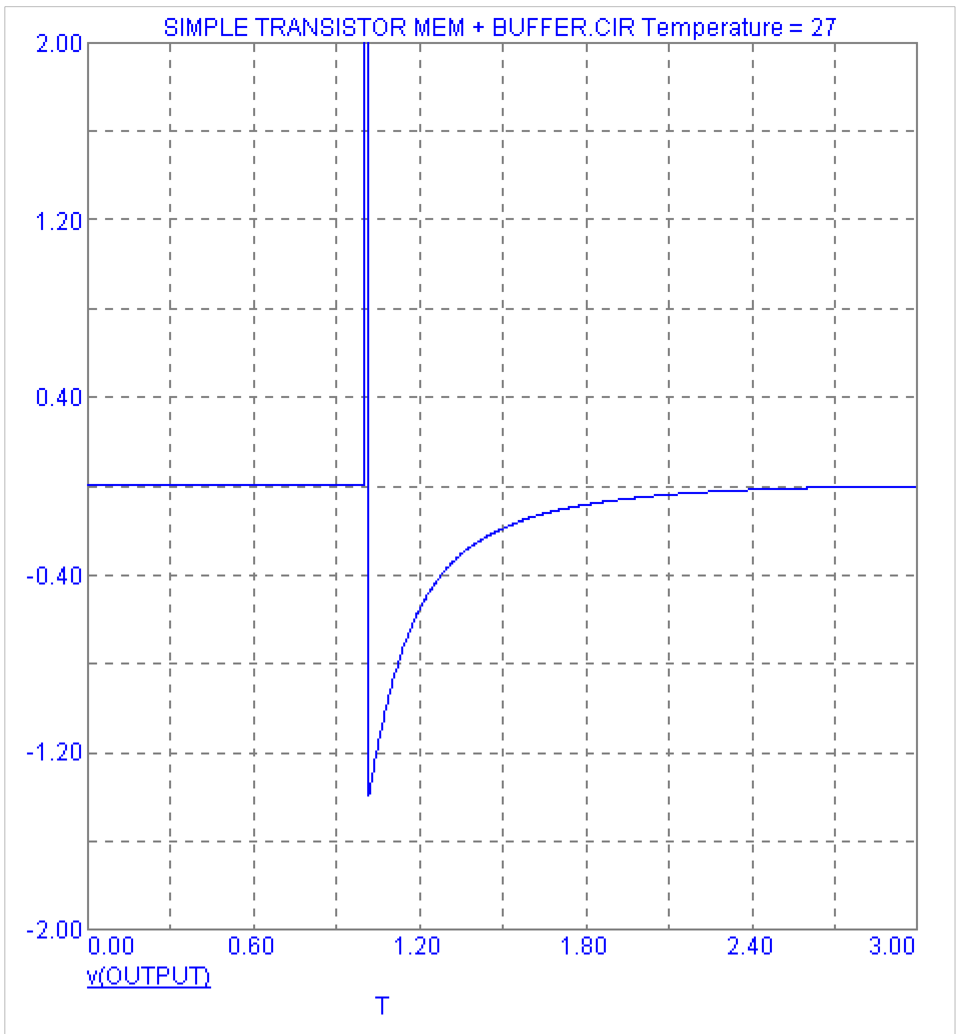

Fig. 1-7b : Output voltage (zoomed)

This transient heats Q1 by 0.215°C ; while this is very small, is is enough to create a noticeable variation of 1.37 V in the output voltage (Fig. 1-7b). This is because the very high input impedance of the buffer magnifies the minute current variations caused by the Vbe delta.

We can try other test signals :

- 1Hz at 2.83VRMS (Fig. 1-8a) not a common frequency per se in music, but on the order of magnitude of the variations of the power envelope of the signal, which influences power dissipation in Q1. On Fig. 1-8a, we see the junction temperature varies of 0.38°C peak-to-peak, which would translate into a 2.42V output drift.
- 10 Hz at the same amplitude, seen in Fig. 1-8b.

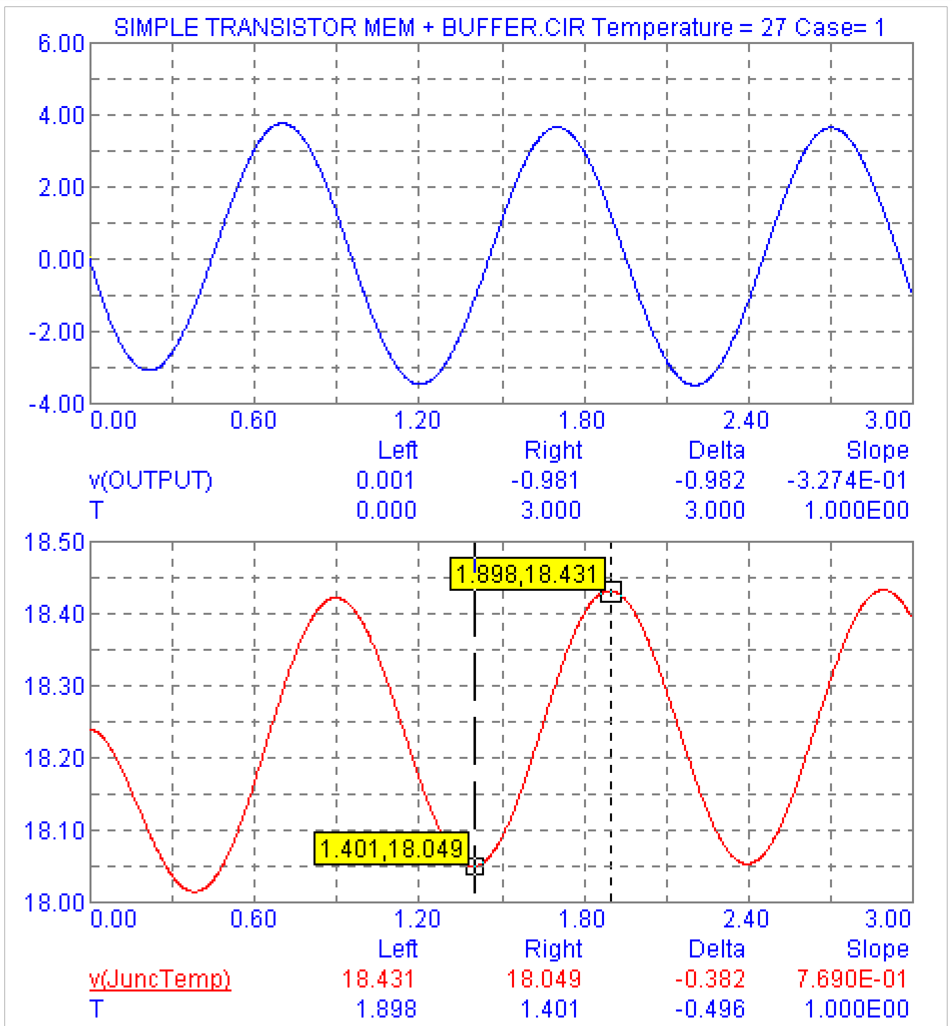

Fig. 1-8a : Temperature variations on a 1 Hz tune

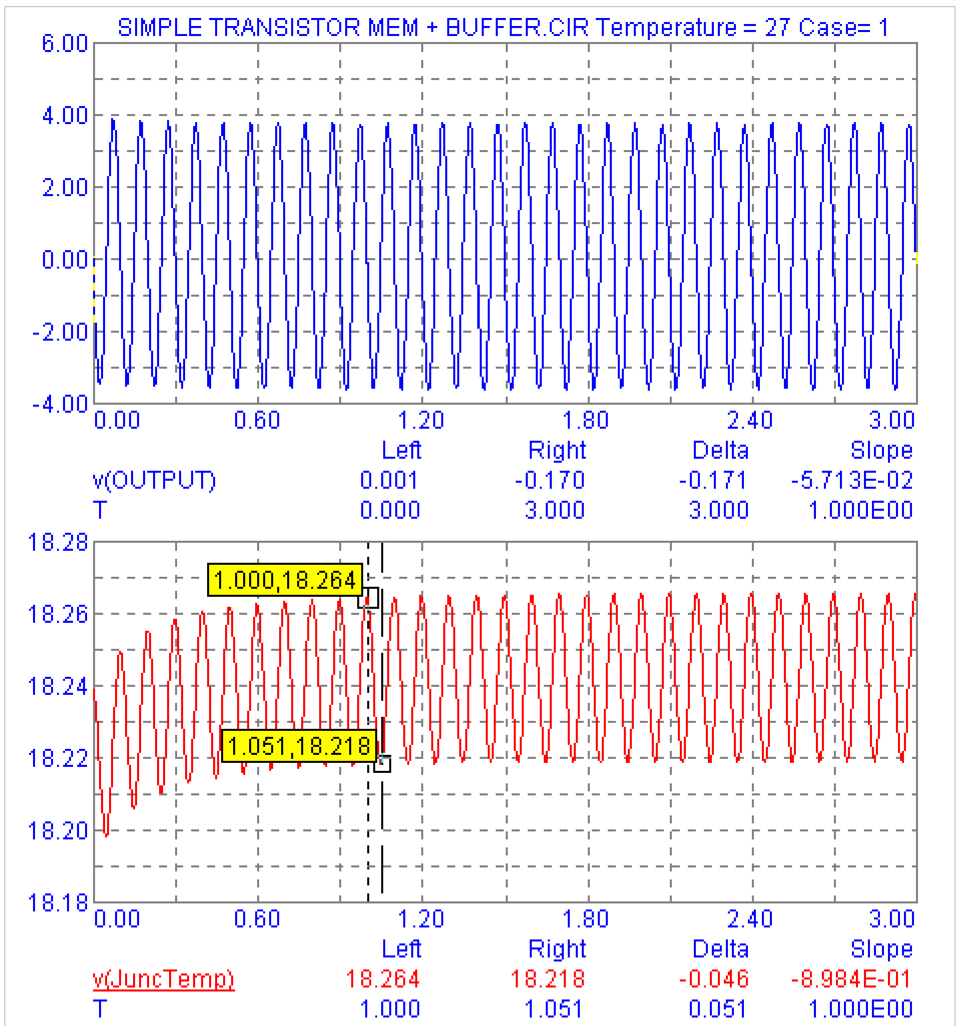

Fig. 1-8b : Temperature variations on a 10 Hz tune

These output drift voltages will not appear at the output of the amp, as they will be corrected by global feedback. However, they will cause the DC bias of the whole amplifier to shift a little as the feedback corrects the DC error. The operating points of the input pair will change as well as the one of the VAS. As we have seen earlier, this will change the distortion spectrum of the amplifier, hence, its sonic character will be dependent on the power of the signal that has passed through it in the last seconds. It will not sound the same after a loud passage than after a quiet one ; powerful Bass notes will upset the balance too. It takes very little to upset an input pair, and Douglas Self measured that an imbalance as little as 1% in current between the two transistors is enough to generate a measurable increase in distortion.

As I wrote earlier, this could be considered like intermodulation, but it is intermodulation between the audio signal and the hidden variations of DC conditions. It is different from simple intermodulation between two frequencies, as this thermal drift signal can not be measured because the feedback corrects it. The drift signal also depends on the original signal, which makes it even worse.

5. The input stage wanders around

Now, consider the model amplifier of Fig. 1-9. It has no output stage, and only the VAS transistor has a thermal model, to keep things simple.

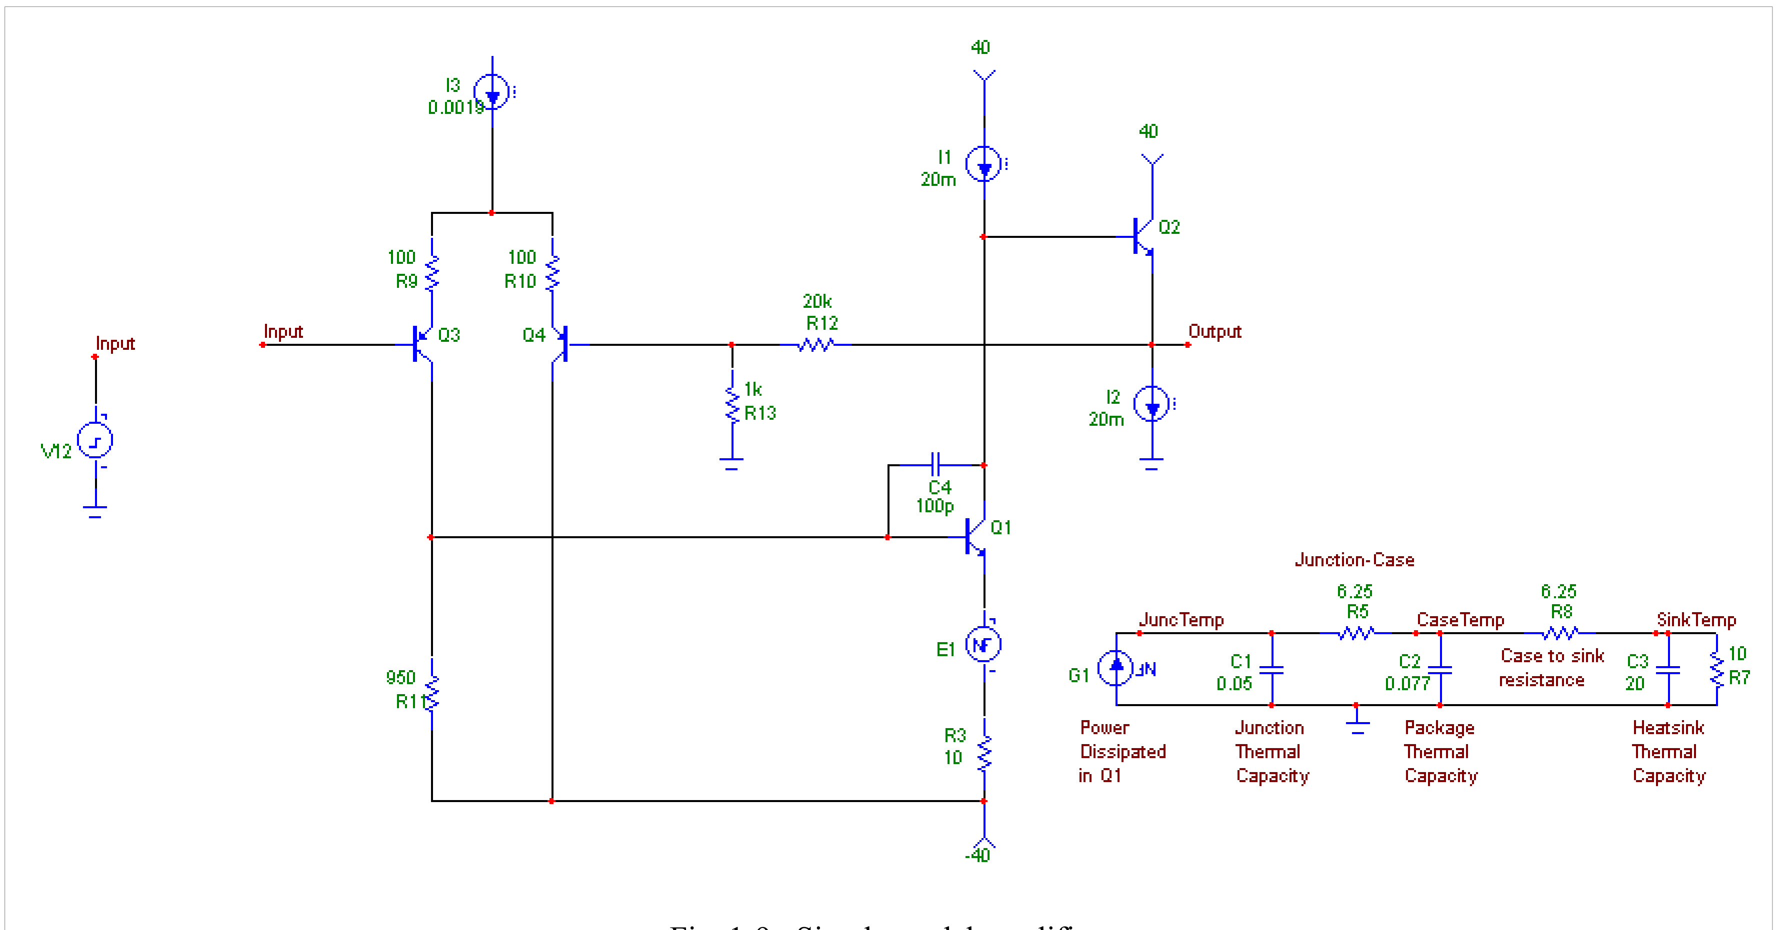

Fig. 1-9 : Simple model amplifier

Having this amplifier process the large transient of Fig. 1-7 gives us the curves on Fig. 1-10 :

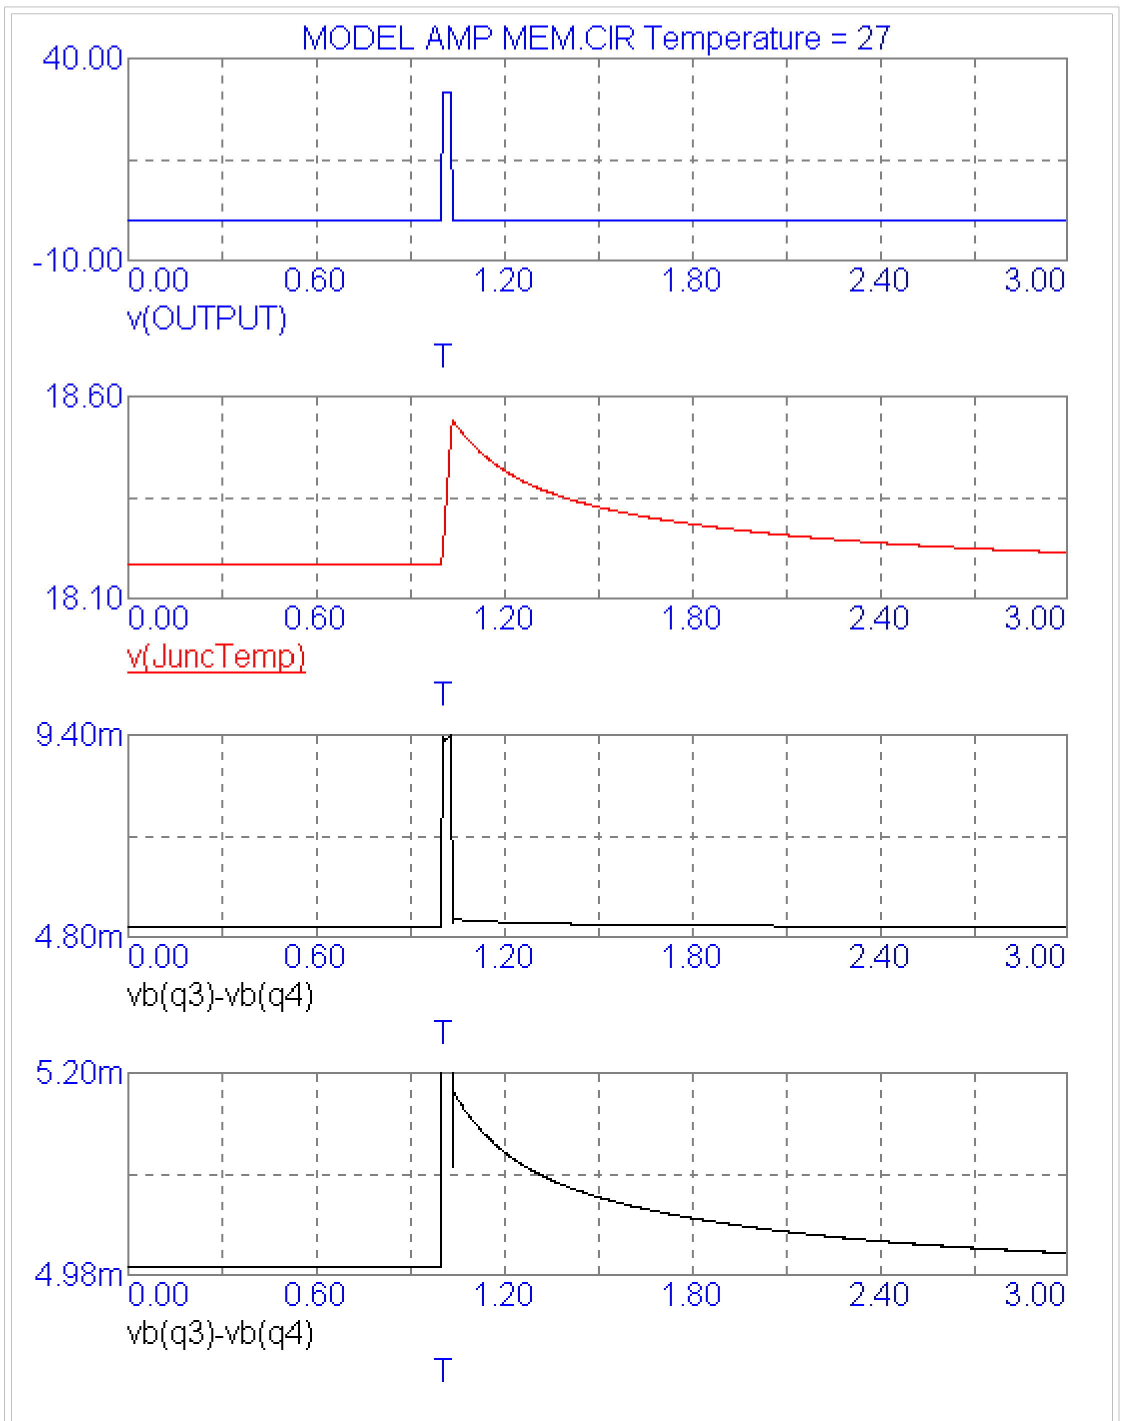

Fig. 1-10 : Thermal drifts in simple model amplifier

Please note the amplifier is not driven into clipping. The temperature delta on Q1 is still there (red curve), but the output voltage (blue curve) shows only a measly 4 mV drift, instead of 1.37V : global feedback did its thing and nicely hides the gremlins from sight. We can still hear them scream, however.

The black curves show the drift of input pair error voltage Vb(Q3)-Vb(Q4) ; the bottom one is a closeup :

- At T=0, VError = 4.98mV
- During the impulse, for an output of 30V, VError = 9.302mV (This gives a rather lowish openloop gain of 6940)
- Immediately after the transient, VError = 5.177mV

The thermal drifts therefore caused VError to shift of 197uV (this corresponds to the 1.37 value at the output if we multiply it by the OL gain). If we compare this to the VError during the impulse (4.322mV), we come to this shocking conclusion :

The thermal drift signal is only 27dB down from the musical signal !

Besides, if we neglect the base current of the buffer transistor Q2 (which is perfectly rational), Intensity in Q1 is constant and dissipated power varies linearly with signal amplitude. Therefore, Q1 thermal drifts are proportional with music signal amplitude, which comes to another bad thought :

This 27dB ratio does not depend on the signal amplitude ! but only on circuit design. It will affect the sound quality of the circuit at all levels.

6. Tubes vs. Transistors

Of course, tubes don't care about temperature. This might be the reason why they sound better...

7. Feedback

Feedback hides the bias variations : as one stage gets imbalanced, the other stages will tend to get imbalanced too, but the opposite way, in a desperate effort to compensate.

We saw above that the thermal drifts of the VAS transistor were multiplied by its gain. Reducing the open-loop gain (by loading the VAS collector to ground with a resistor) therefore reduces the drifts, but it also reduces the feedback factor available to correct them, resulting in a draw. Besides, it adds distortion.

[Flame suit on] I personnally think that high feedback is preferable. TIM seems more like a legend to me. High feedback can be used more efficiently against memory effects. Besides, you can make a fast amp that has high feedback too. There will be more on this in the next article, which will cover the details of the amplifier I designed against memory distortion. **[Flame suit off]**

Further reading

- There is a very interesting article that shows how to actually *measure* these elusive thermal drifts (hint : it is based on the integrated dc error over time). [Download]
- I spared you the maths, but you can find them in all their glory in the patent on the global amplifier design. Recommended reading. [Download]
- The patents on the Lavardin input stage : interesting, bur read the other ones before that one. Besides, my input stage is funkier. [Download]

All files ate zipped. Images are in TIFF CCITT4 Format, which allows for excellent compression ratios (1/4 the size of Gif's). You can read them with ACDSee, Photoshop, Paint Shop Pro, and the like.

I'll be back

I was into electronics since a long time, but mostly digital. I had no idea how an audio amplifier worked until september, 2000 (!). If you want to know how I designed a small amplifier that looks impossible on paper, yet makes my reviewer-recommended, purist, prized-by-La nouvelle revue du son, $800 store-bought mosfet amp sound like a pocket AM radio, don't miss the next article !!!

Memory Distortion - Part 2 : The Input Stage

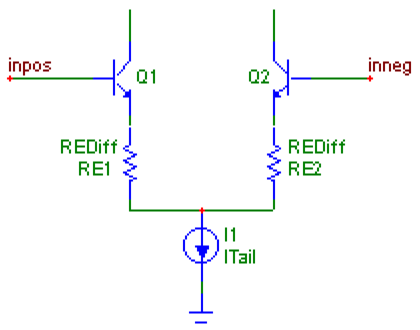

There are billions of these critters around the world : think about all the opamps in your TV, CD player, portable phone, etc. This circuit topology has become so common that almost everybody takes it for granted. With the very notable exception of Tubes, audio amplifiers almost universally use a differential pair. Those which don't have a differential are often exotic designs, like zero-feedback, single-ended class A oddities.

Résumé of previous épisodes

If you read the previous article, you remember that the *Memory distortion theory* states that thermal pulses in the transistors, caused by the audio signal, upset the global DC operating point. Feedback hides the bias variations : as one stage gets imbalanced, the other stages will tend to get imbalanced too, but the opposite way, in a desperate effort to compensate. Therefore, the distortion spectrum, which is dependent on the operating point, varies according to the signal which passed through the amplifier in the last few seconds. Tubes are insensitive to this phenomena, which could be the explanation of why transistors and tubes sound "different".

Input Stage Simulation

Fig. 2-1 is the schematic, in which the two voltage sources taken in the emitters represent the Vbe shifts with temperatures. The output is taken as I(R1)-I(R2). The input, or *error signal*, is the voltage between the two bases. **RE**=100R, **ITail**=8 mA.

Fig. 2-2 shows the dissipated power in both transistors versus error signal (which is simply the differential signal betwen the two bases). The high rail voltage makes them dissipate quite a lot ; the maximum value (at clipping) is around 250 mW. I modified the thermal model for small TO92 transistors, BC550 for instance, which are much easier to heat up than the bigger ones we use as VAS.

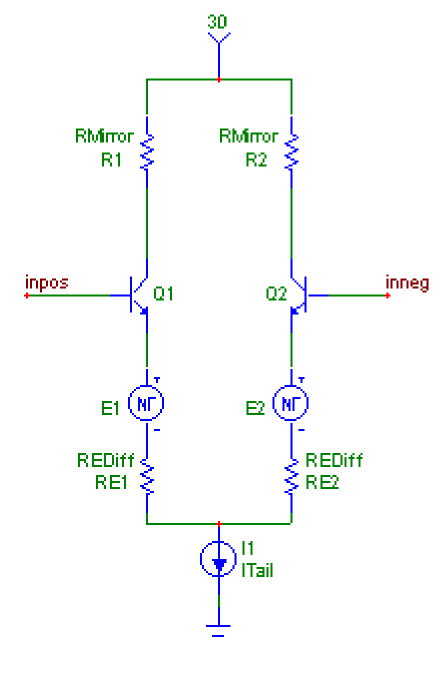

Fig.1 : The differential pair.

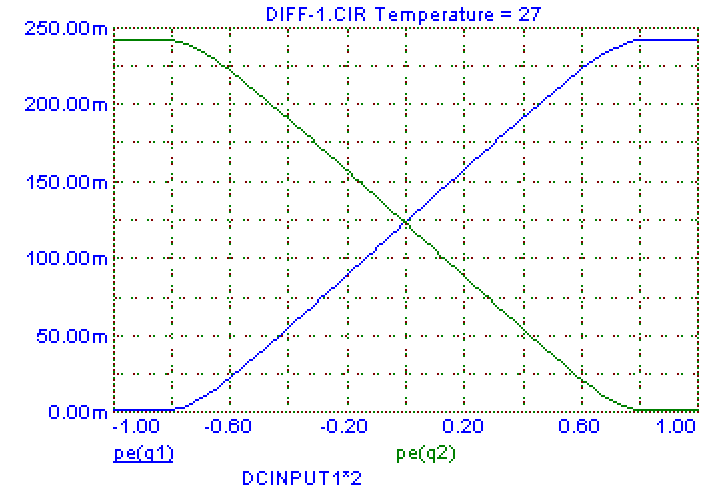

Fig.2 : Dissipated power in Q1 and Q2 vs error signal.

Consider an amplifier with the following parameters as an example :

- **G** : closed-loop gain (20 dB, or 10)
- **Ao** : openloop gain (46 dB, or 200)
- **FB** : feedback factor **Ao**/**G** = 26dB, or 20

We will suppose for a moment that we have an amplifier with large openloop bandwidth, in which **FB** is constant with frequency, to simplify. Now, let's suppose this amp is asked to produce the same transient as in the last article (30V for 20ms) : **Vout** = 30V at the ouput corresponds to **Vout**/**Ao** = 150 mV between the transistor bases.

In Fig. 2-3, I plot the die temperature of the two transistors. The black curve is the output voltage, the green and blue are the temperature of both transistors.

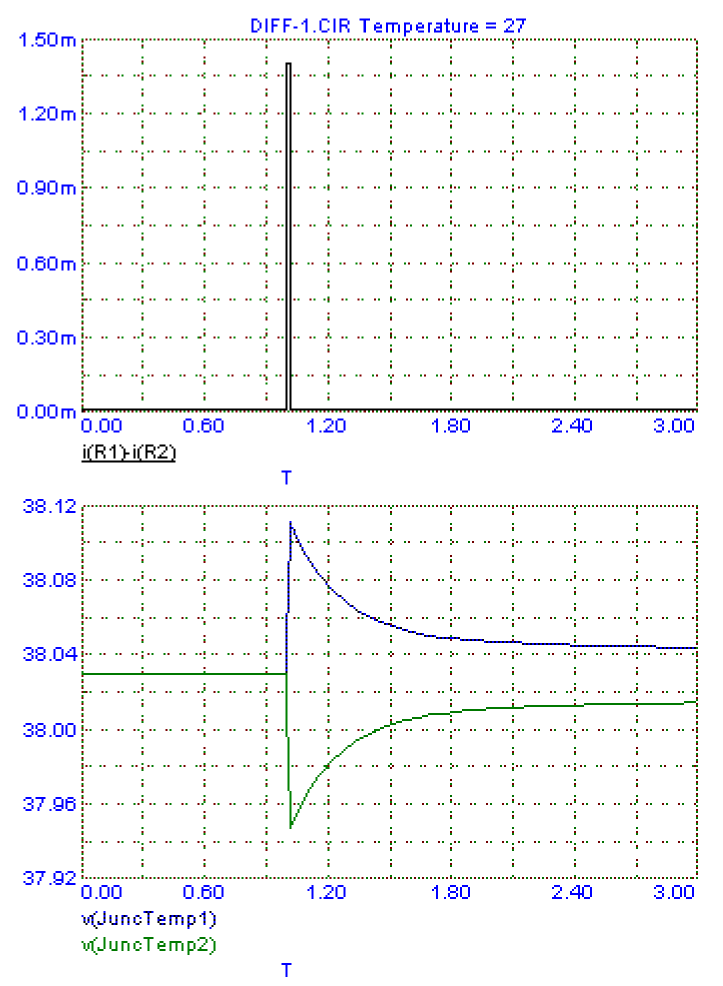

Fig. 2-3a : Temperature of Q1 and Q2 versus Time

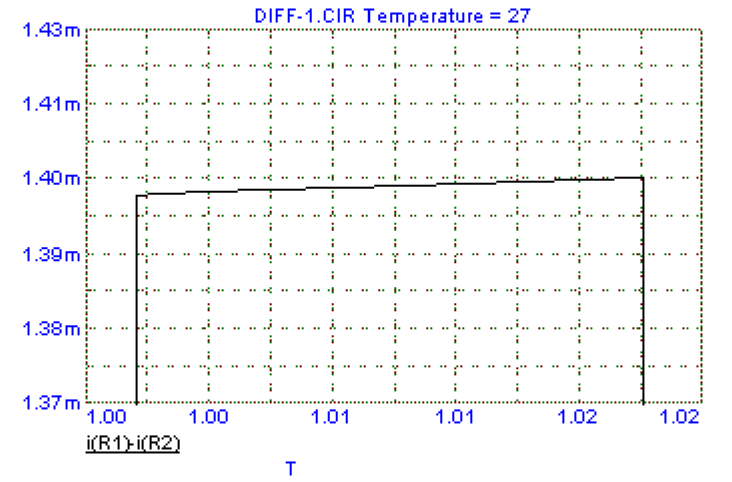

Fig. 2-3b : Output Signal (zoomed)

In this example, the transient generated a temperature delta of 0.016°C, which is pretty small. It translates as a 3uA difference in the output currents, which is **50 dB** smaller than the audio signal. This thermal signal is quite small. However, these thermal drifts have importance regarding sound quality, as we will se later.

Before I forget, let me mention the interesting fact that, no matter what values we have for **FB** and **Ao**, this **50dB** between thermal signal and error signal is likely to be a constant, as thermal dissipation is proportional to error signal (see Fig. 2-2). So, we have not found an explanation on why the low-feedback amps seem to sound better.

Forgot something ?

Of course, we forgot the audio signal itself. The error signal is differential, but the audio signal, which comes at the input of the amplifier, is common mode. Fig. 2-4 shows what happens when both transistors have the same base voltages : the purple curve is their dissipated power vs. common mode voltage, and the green one is their junction temperatures, while processing the same transient as before.

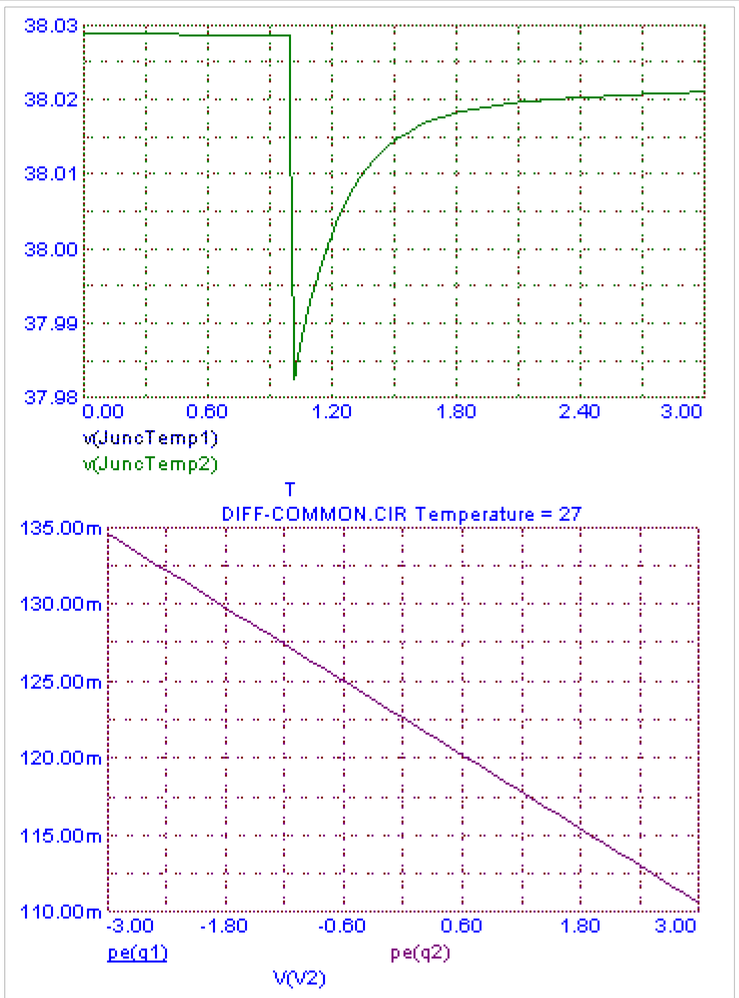

Fig. 2-4 : Thermal variations in common mode

Their temperature varies of 0.046°C on the transient. This is much higher than the variation caused by the error signal.

On Clipping

Much has been said on clipping. A fact is that amplifier behaviour at clipping seems to be an important factor for sound quality.

At clipping, feedback no longer maintains a low error signal : the error signal is simply the difference between the actual Input Voltage and the maximum input voltage that can be amplified without clipping. The same transient as before, driven into clipping with a 1 Volt error signal, generates a delta of 0.1°C between Q1 and Q2.

Questions, questions

Unlike in the last article, where the numbers clearly showed that there was something happening, this time the thermal effects seem rather small. Effects caused by the error signal really seem negligible. The other two, common mode audio signal and clipping, might not. Let's see what happens when the openloop gain changes.

In a low feedback amp (**Ao**=200 or 46dB), the error signal for a full scale output (**Vout**=30V) is around **Vout**/**Ao** = 150mV. A delta of 0.1°C, offsets the error signal by 200uV, which is very little indeed.

In a high feedback amp (**Ao**=50000 or 90dB), the error signal will be 600uV for a full scale output. Then, 200uV generated by one tenth of a degree of thermal imablance do not appear to be negligible anymore. It is actually one third of the actual audio signal !

High-feedback amps have a low open-loop bandwidth, as low as 10-100Hz. This is associated with the notion of them being "slow", i.e. incapable of reproducing transients, while in fact they have very good speed and slew rate. The reason why they often sound "slow" on transients might be that, when they clip, the tiny thermal imbalance that is generated will be enough to offset the operating point of the input stage, and then the other stages, by a noticeable margin, for several seconds after the instant when clipping has occured. The quasi-religious question of feedback is not solved, but we know a little more.

The clipping behaviour of tubes vs transistors, has been compared many times, especially in amplifiers which have to deal with original signals coming from microphones, which have much more dynamics than recorded ones. However no really conclusive evidence of why tubes clip more gracefully has appeared. This thermal imbalance theory might be the answer to that old question...

Finally, it has appeared in this page that the input stage did not generate much thermal drift (except on clipping). It is, however, sensitive to the drifts generated in the other parts of the amplifier, especially the VAS (see last article). We will take this into account in the input stage design.

Memory Distortion - Part 3 : Measurements

How to measure Memory

The method that Gérard Perrot, author of the patents, used is very clever. The drifts we want to measure are hidden by feedback : as one stage gets imbalanced, the other stages will tend to get imbalanced too, but the opposite way, in a desperate effort to compensate. This must result in an extremely small drift voltage at the amplifier output, which is dependent on the actual internal drift, but much smaller. In order to detect it, integration of this offset over a period of time is used, as explained in the AES article that you can download by following this link.

Example of measurements

Using the method described in the article, I simulated a number of circuits and their memory measurements. The test waveform was to make the amp output 20V during one second, and the thermal models used were pessimistic, to exacerbate memory effects and ease the search for a memory-free circuit topology. I started with the circuit of Fig. 3-4, and attached a thermal model to all transistors.

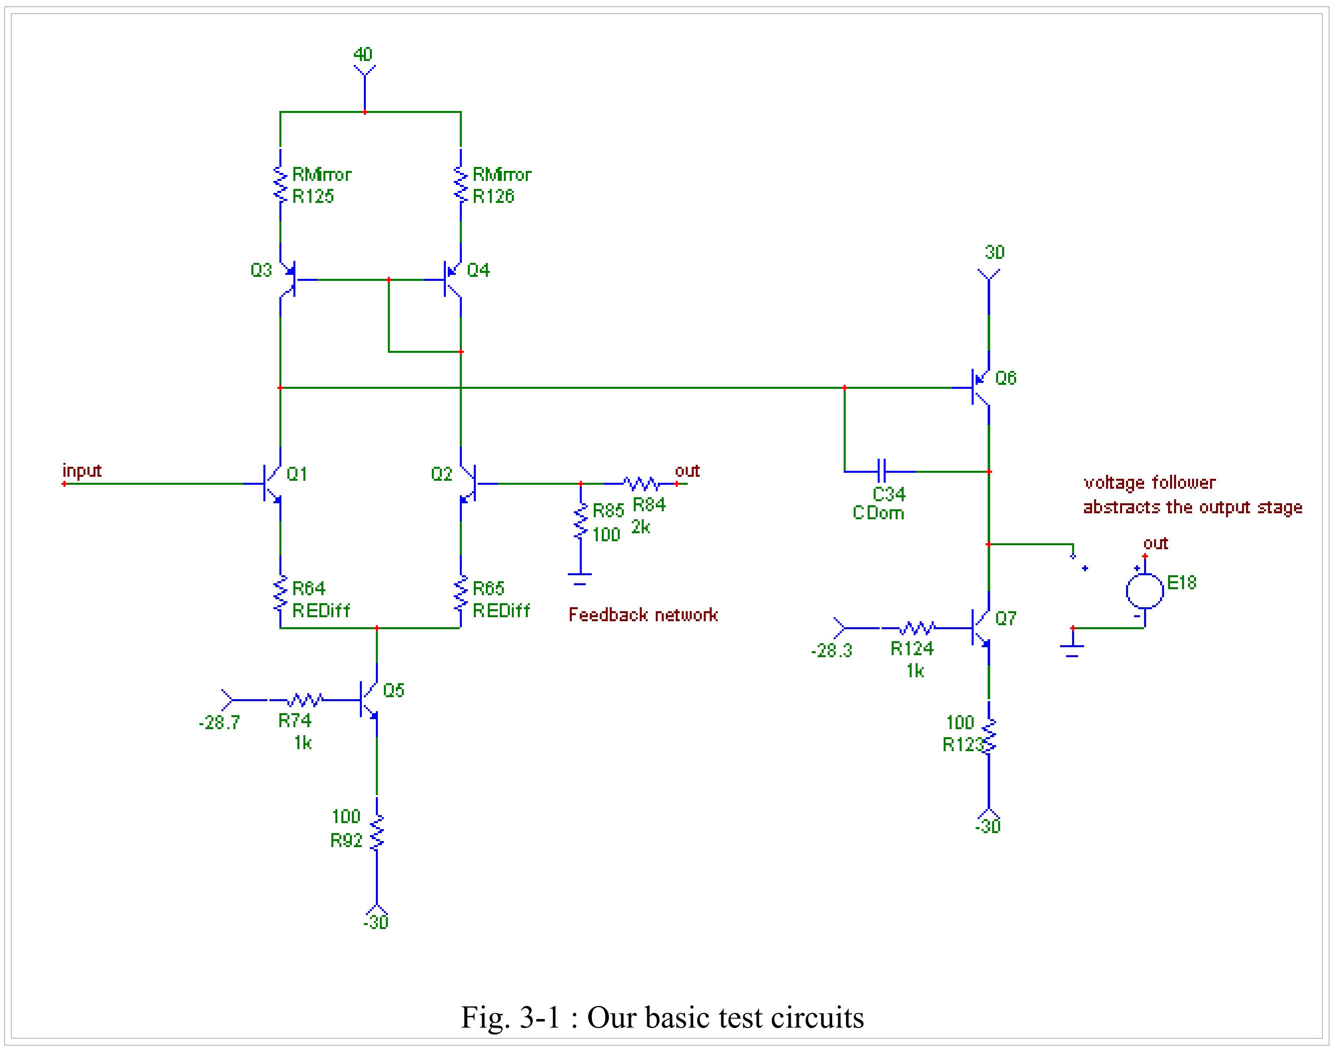

Fig. 3-1 : Our basic test circuits

The resistors in the current sources are 100 Ohms, and the other circuit values are :

ReDiff	100 Ohms
RMirror	100 Ohms

I will elaborate on the first measurement of Fig. 3-5 : you will see more of these strange-looking curves. This curve is simply the integration (sum over time) of the output offset voltage. Integration starts right after the end of the test waveform (we don't want to measure the waveform itself, but the memory of it). Once the test waveform is over (at t=1s), the output should return to zero volt. It doesn't however, because the bias points have shifted. Instead, the output voltage of the amp slowly comes back to normal as the transistors come back to their usual temperature. Integration shows us the total error summed over time, and this is a practical tool for evaluating the total memory of the amp.

This gives us the curves in Fig. 3-2 : blue is output voltage, black is integrated memory signal. We can see the tiny variations of the output voltage in the zoomed curve on the right. Remember, these seem very small because they are hidden by feedback.

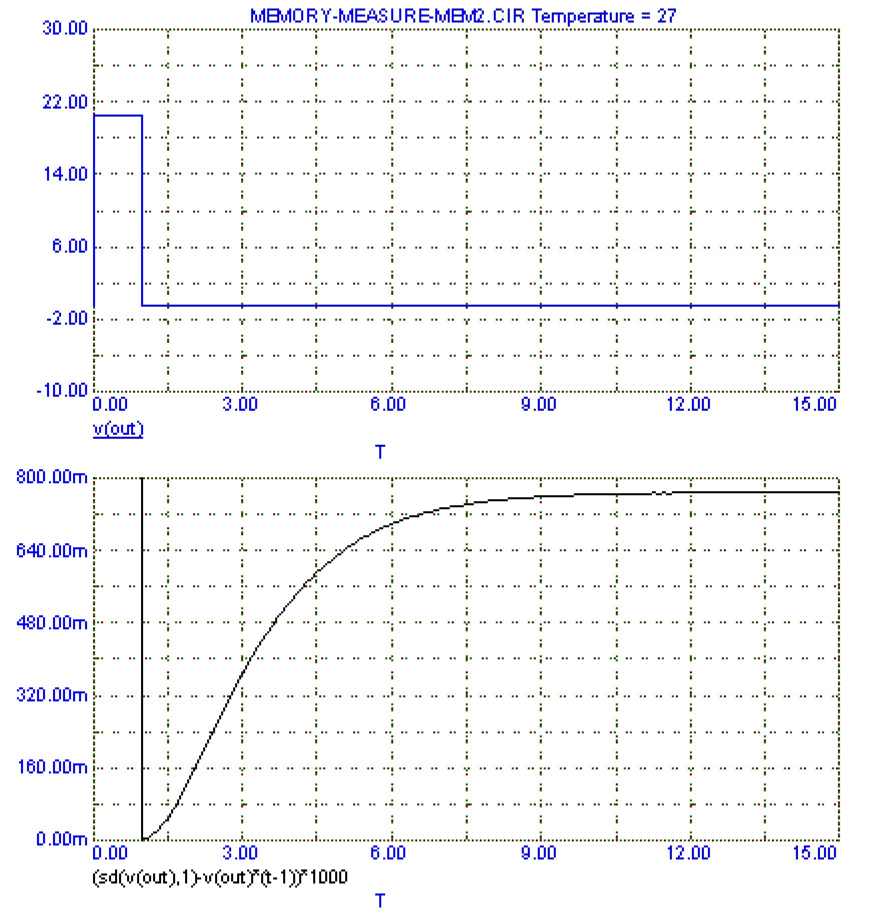

Fig. 3-2a : Memory, input stage only

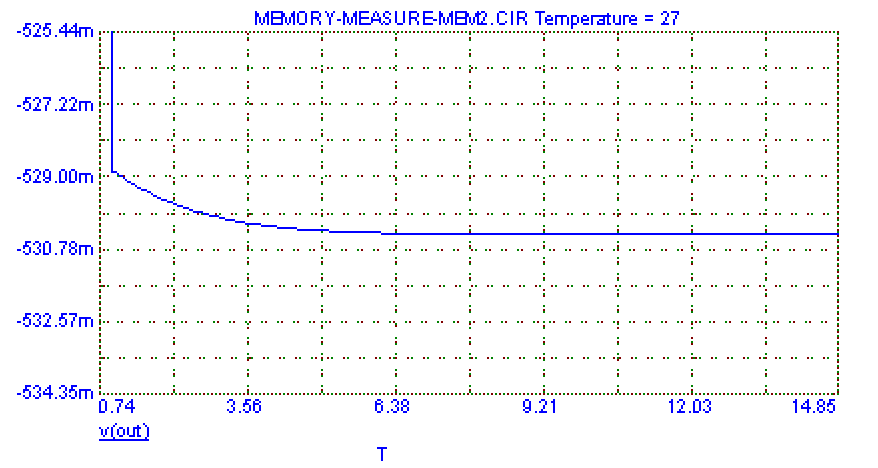

Fig. 3-2b : Output voltage, zoomed

We will use the value the curve reaches at the far right (here 2.66) as a measurement of memory. These are arbitrary units, so the only use of such a value is for comparison. So, I will activate the thermal models in all the stages of the amp and plot their respective memories. I will not include all the curves here (they all look the same), but only their final values.

Of notable importance is, also, that I used the VAS transistor as a Current amplifier, whereas in article 1 it was used as a Voltage amplifier. In the voltage amplifier configuration, with resistors instead of the current mirror, Q6 generates much more memory, because the thermal Vbe errors are multiplied by the VAS gain, which is large. This gives memory measurements an order of magnitude higher than with the current mirror. Q6 Beta varies with temperature, also, but apparently this is much less problematic. It seems we will stick with the current mirror, then.

Thermal model activated in...	Memory
None.	0.00001 (roundoff errors)
Input pair (Q1,Q2)	0.77
Current mirror (Q3,Q4)	0.27
Tail current source (Q5)	0.3
VAS (Q6)	4.76
VAS Current source (Q7)	2.7
Both current sources	3.
All at once	4.83

Apparently we were a little bit over optimistic in part 2 by saying that the input stage would not really create offset on its own. It does, and quite happily, even if it is not the main source. These "almost negligible" values might not have been so negligible after all...

We will now apply the thermal models to all transistors except one stage, to see in which stages eliminating memory will benefit the most.

Thermal model activated everywhere except in...	Memory	Gain
Input pair (Q1,Q2)	3.08	1.75
Current mirror (Q3,Q4)	10.6	Ouch !
Tail current source (Q5)	3.72	1.11
VAS (Q6)	0.63	4.2
VAS Current source (Q7)	3.86	0.97
Both current sources	2.68	2.17

The Current mirror looks suspicious ! There must be an interaction between this and other stages which partly cancels with all thermal models on, but deactivating the thermal model in the current mirror prevents this cancellation and gives strange measurements. I guess the only solution will be to optimize all parts of the circuit regarding memory.

In the next chapter, we will examine each of these circuit elements and see how we can get rid of their memory by adapting their topology.

Memory Distortion - Part 4 : Circuits

On to the exciting stuff !

The Constant Power concept

Consider a transistor. If we neglect its base current, the thermal power it dissipates is :

p = **Vce * Ie**

If we wish to reduce thermal effects, we want to enforce a constant junction temperature, hence a constant dissipated power, in all critical transistors. This can be done in two ways :

- **Vce** and **Ie** are constant
- **Vce** = Constant * 1/**Ie**, or an approximation of this.

Now, on to the applications.

Current Sources

They are simpler to explain and will make a good example before we start with the other, hairier stuff. Conventional wisdom has it that "current sources sound bad". Maybe this is because they add memory, as the measurements showed.

The critical transistor of our current source (Fig. 4-1) is Q1. Thermal Vbe variations will change the voltage on the emmiter resistor, hence the current.

The nice thing in a current source is that, obviously, the current is almost constant. Therefore, we will just put a cascode transistor (Q3) to enforce constant **Vce** on Q2, and then our constant **Vce** and **Ie** is done !

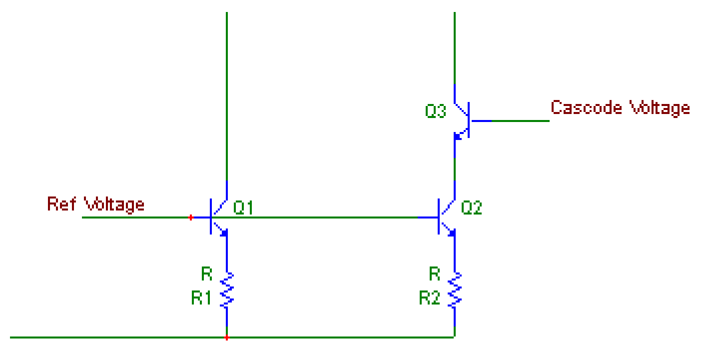

Fig. 4-1 : Current sources

Here, I activated the thermal model only on the current source transistors (and the cascodes) of both sources (input pair tail and VAS), but not in the rest of the amp. The measurements are self-explanatory :

Circuit	Memory
Simple source	3.
Cascoded	0.016
... and Re=1K	0.004

It was known that cascodes improved the sound, but it was unclear wether or not it was because they give better frequency response. Doc Bottlehead sells cascode current sources to put in their highly famous single ended tube amps, so they must sound good. These results show that the effect of cascodes might be simply thermal.

From now we will keep the cascoded current sources, with **Re**=1K in the input stage and **Re**=650Ohms in the VAS. Voltage references will be taken from 8.2V Zeners, which is the voltage for which the Zeners are the less noisy.

Now that we have seen what tweaking an innocent-looking current source can do, let's tackle the input stage.

The input stage

In order to reduce memory, we must ensure Constant thermal dissipation in critical transistors, regardless of :

- Common mode input voltage
- Differential input voltage
- Tail current (more or less)

Besides, this input stage should be insensitive to the drifts generated by the other stages, especially the VAS.

I will spare you the details, but I tried countless configurations, more or less inspired by the Lavardin patent [Download]. Their solution (read the patent !) is elegant but there is a hint of positive feedback, and it does not take into account the differential voltage. This is not very important as their amp is a very high feedback design, and error voltage is very small, but still, it leaves room for more brain-itching.

I finally arrived at the following configurations :

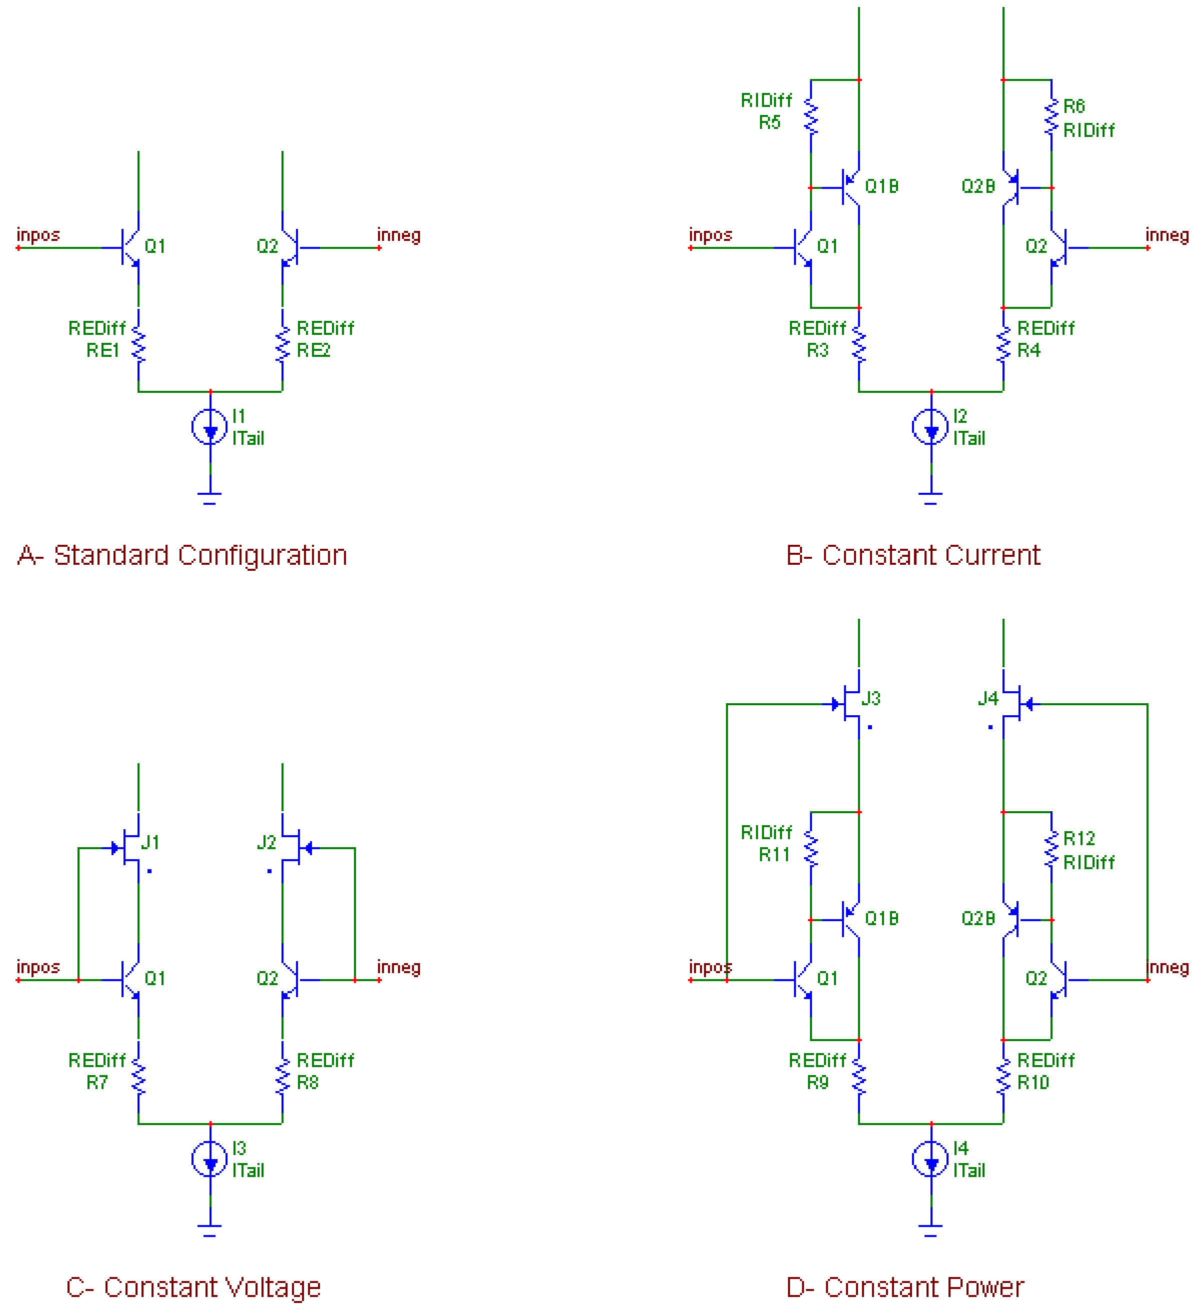

Fig. 4-2 : Input Stages

Configuration B is a **Complimentary Feedback Pair**. The current in Q1 and Q2 is simply the **Vce** of Q1B and Q2B divided by RIDiff. This **Vce** is quite constant, so the current is, too. However, Q1/Q2 **Vce** is not constant as it is simply one **Vbe** lower than the **Vce** in the standard case. Thermal power is not constant but should be lower, because the current is lower.

Configuration C is a simple **cascode** with JFETS. 2SK246 will suit this application well (go to Mr Borbely's site and read the articles on JFETs and cascodes). It will maintain its source around 2 volts above its gate, which is perfect here, as it will give us a constant **Vce** of around 2.6V. Mr. Borbely connects the gate of the JFET to the emitter of the transistor below in order to neutralize its parasitic gate-drain capacitance. It could work here, with one transistor.

Configuration D is the combination of B and C : a Complimentary pair to enforce constant current in Q1/2, and a Cascode to enforce constant Vce. Try to analyze the circuit, you will see it works. Q1 and Q2 are operated at **constant power**. Neat, eh ?

Of course, it looks complex and ugly, but it is the simplest way that I know of to achieve constant power. As an added bonus, the local feedback of the complimentary pair makes it **much** more linear than the standard configuration. And, the parasitic G-D cap on the JFETS prevents the whole mess from oscillating : it will **not** do dirty stuff at high frequencies.

Before measuring memory, let's make sure this thing is fit for the job by plotting the transfer function of Output (diff current) versus Input (Diff Voltage) (see Fig. 4-3).

The red curve is the new input stage, the black one is the simple, two-transistor version. Top curve shows Iout versus Vin, and the two curves at the bottom show the gain variations (derivative) of the previous curve. Same vertical scale is used on both, so we can see the new stage has a much flatter curve, meaning more linearity. The blue curve is the configuration B, the CFP, which has the same linearity as the new stage.

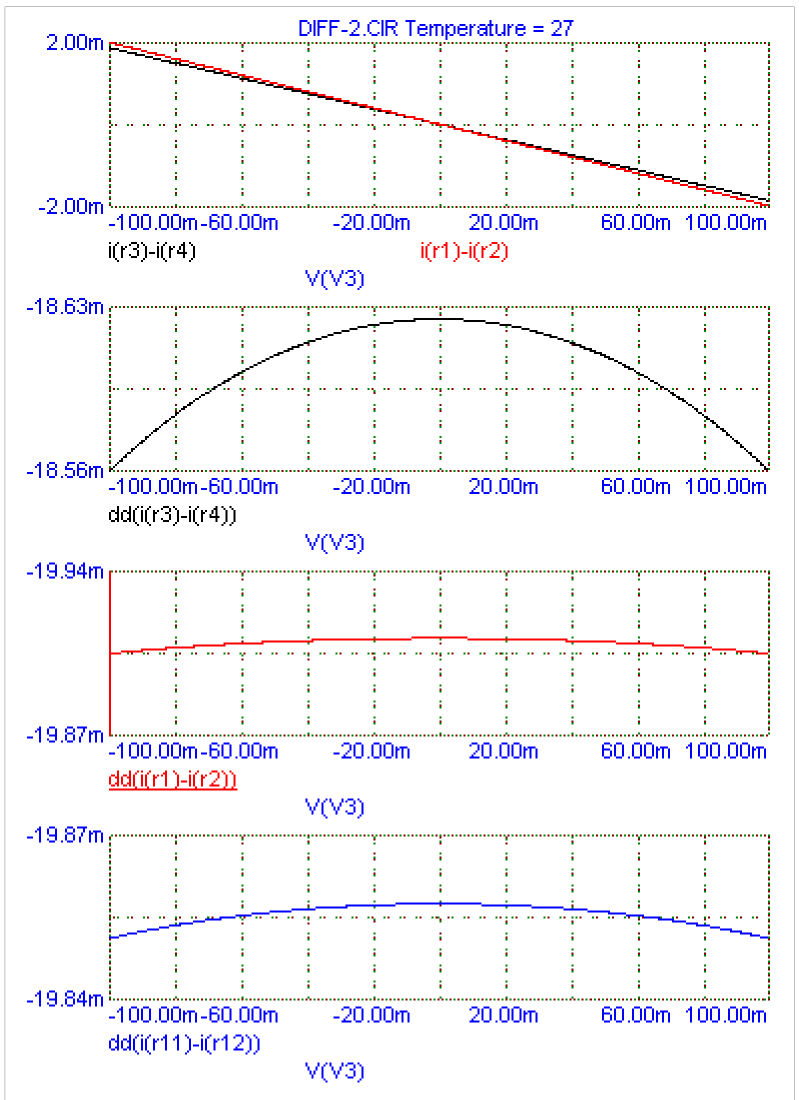

Fig. 4-3 : Input Stage Linearity

It will also be interesting to plot thermal power in the critical transistors Q1 and Q2 versus common mode (Fig. 4-4) and differential (Fig 4-5) input voltage :

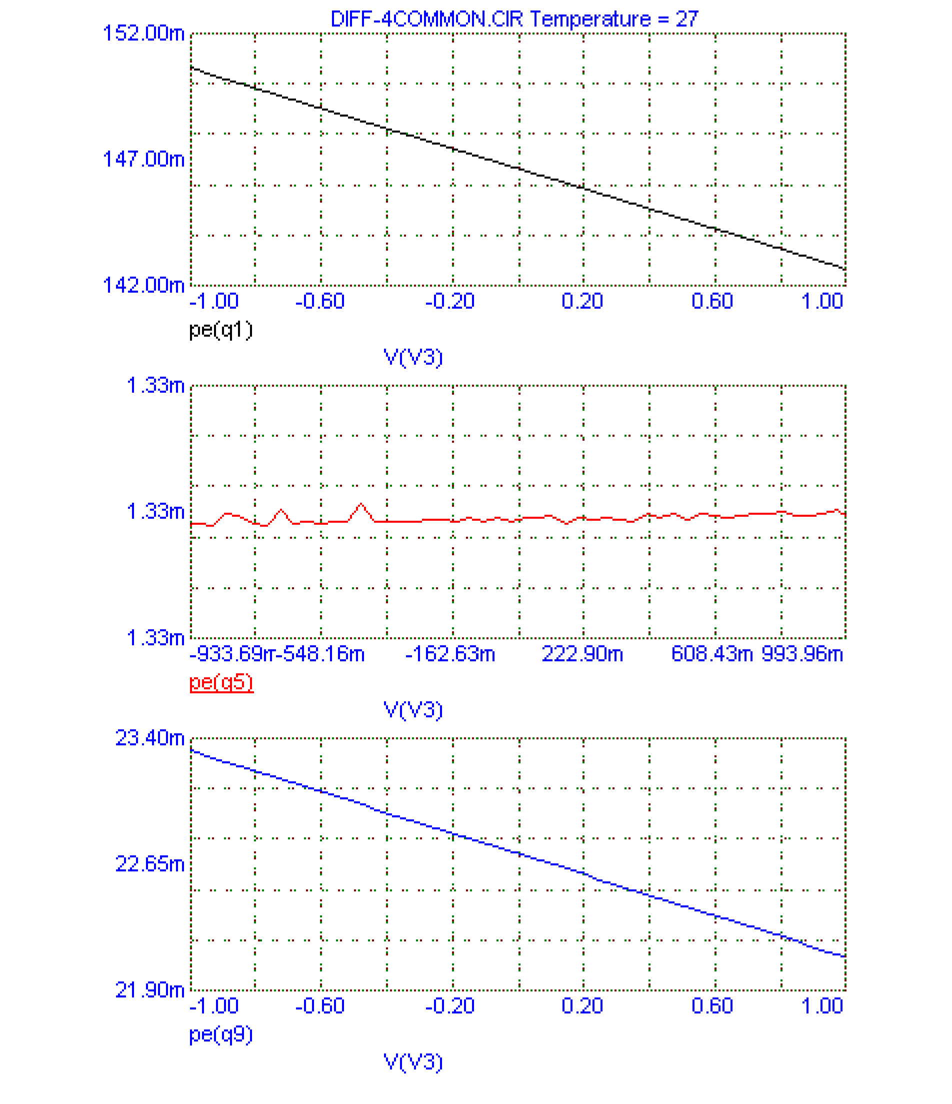

Fig. 4-4 : Dissipated power in critical input transistor versus common mode input voltage

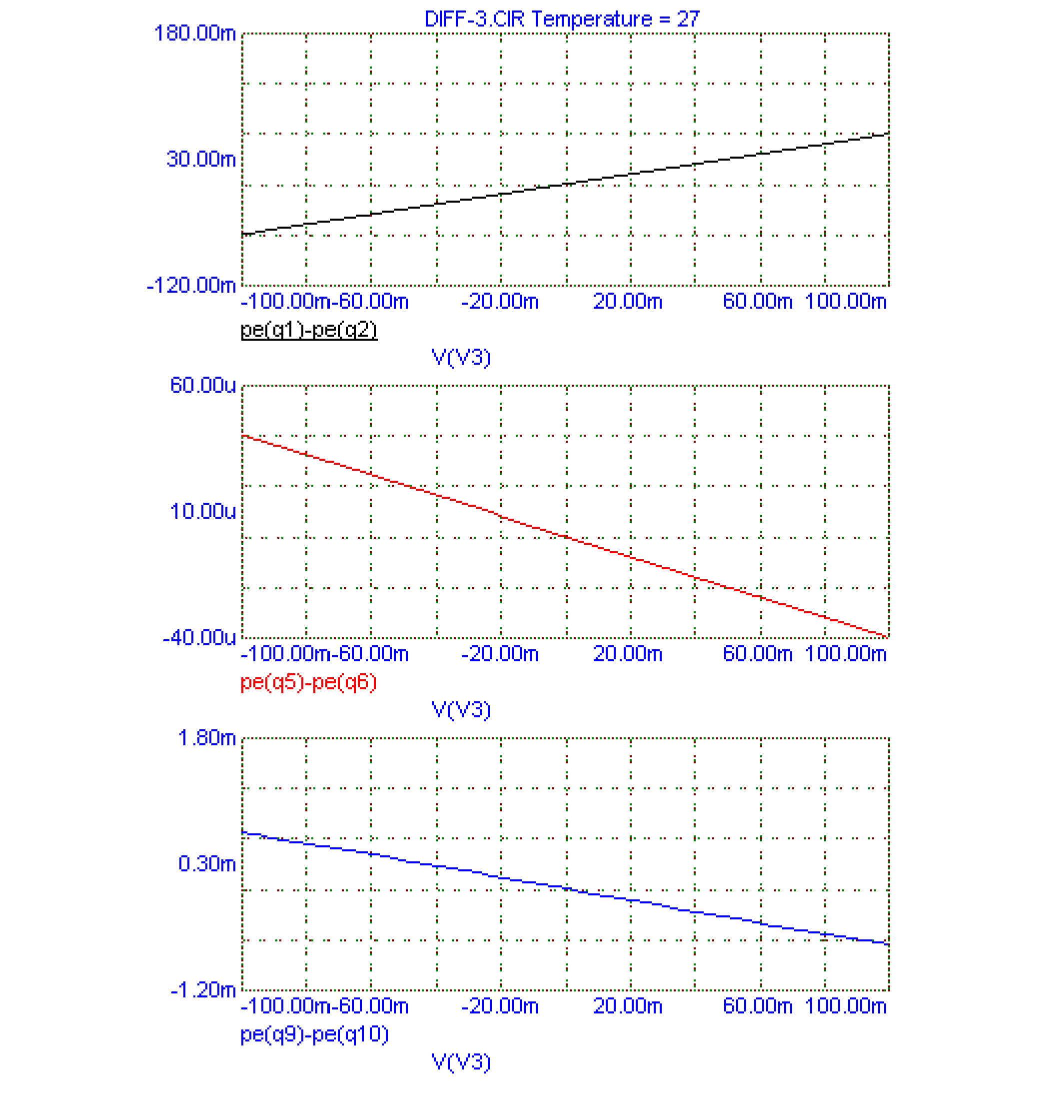

Fig. 4-5 : Delta of Dissipated power between critical input transistors versus differential input voltage

Numeric results are easier to handle (I added the cascode only case). The Common column is the variation in the dissipated power in **one** transistor for a 2 volts variation in common mode input voltage (-1 to +1 volts). The Diff column is the difference in power **between** the two transistors for a differential input voltage of 0.2V (-0.1V on one base, +0.1 on the other). The stages are loaded with 1k resistors int he collectors connected to a V+ of 40V. Powers are in milliwatt :

Case	**Common**	**Differential**
A. Classic	8 mW	60 mW
B. CFP	< 0.001 mW	0.67 mW
C. Cascode	4.7 mW	4.7 mW
D. New	< 0.001 mW	0.04 mW

This is very logical : the cascode and New stages are well protected against common mode. The CFP is well protected against differential voltage. The New circuit is the only one that can handle both with minimal thermal drift. Let's do a "real-life" test by incorporating it in the complete simulated amplifier. Here, the thermal models are active only in the input stage (left column), or in the whole amplifier (right column).

Case	**Memory**	**Memory**

	Input only	Whole amplifier
A. Classic	1.05	3.23
B. CFP	0.293	2.7
C. Cascode	0.192	2.
D. New	0.009	2.

Results are conclusive : it works as advertised. We have to minimize the memory in the rest of the amp, though, it we want to benefit from it. Let's make a quick test to find who the culprit is : current mirror or VAS ? I will also keep the four input stages a for this one.

Case	**Memory** With a perfect VAS	**Memory** With a perfect Current Mirror
A. Classic	2.8	8.
B. CFP	0.276	6.
C. Cascode	0.356	4.68
D. New	0.127	4.56

We have interaction, or compensation effects, again between the current mirror and the VAS, as the perfect current mirror actually gives worse results than the "real" one. Anyway, it is pretty clear from these results, that :

- The New configuration is by far the best one regarding memory,
- The Cascode is better than nothing,
- The CFP alone is not worth the trouble,
- The VAS is the next to get under the microscope.

Memory Distortion - Part 5 : Brainwashing the VAS

It has been remarked that the VAS is important in amplifier sound quality. This backs the memory theory, as its contribution to memory effects is huge. How can a single transistor be so critical ? Simply, it is at the same time the main gain stage and the one that heats the most (it takes the full output swing). And, its errors are multiplied by its gain.

First, rather standard, proposal

For once, I will not have a funky circuit here, and we will simply try to put a Cascode on the VAS.

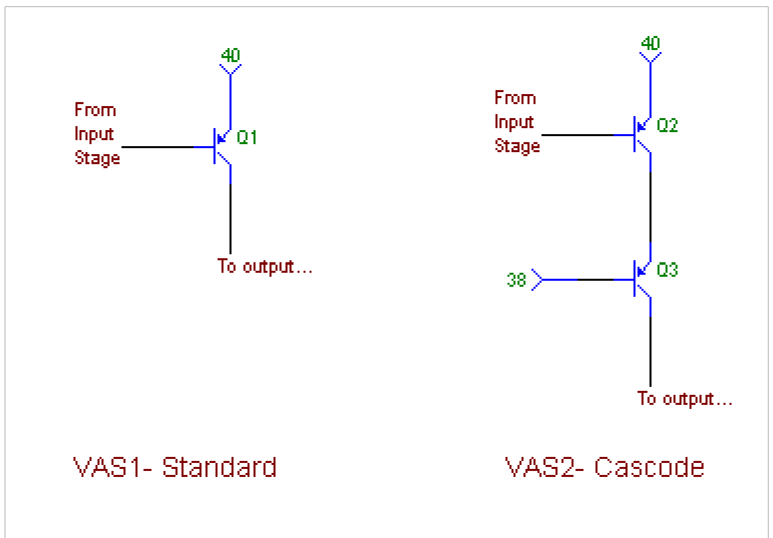

Fig. 5-1 : VAS Variations

We will build a model amplifier around this VAS, make it output a full swing of -30/+30V, and plot the linearity (upper plot) and dissipated power in the gain transistor (lower plot). Blue is for Standard and red is for Cascode.

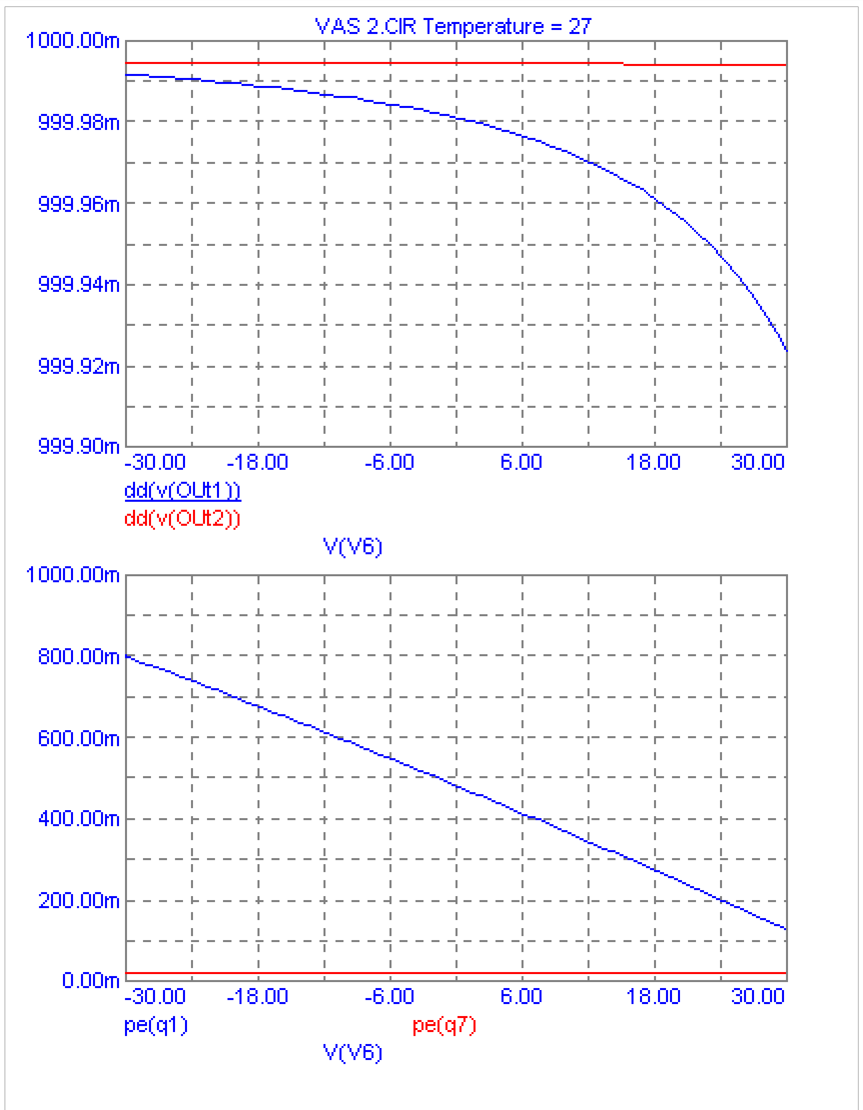

Fig. 5-2 : VAS Linearity and Power

This is old news, but the Cascoded version is much more linear (the gain curve is flatter). See Douglas Self's *Power Amplifier Design Handbook* for details. As for the power, of course the Cascode version has much more constant dissipation. The current does not vary very much actually, as the impedance that the VAS has to feed is quite high, being only the current source collector, and a buffer transistor's base.

Cascoded VAS's are used almost everywhere, and they are a good step towards reducing memory. However, it might be possible to do a even a little bit better.

Enter the Magic Resistor.

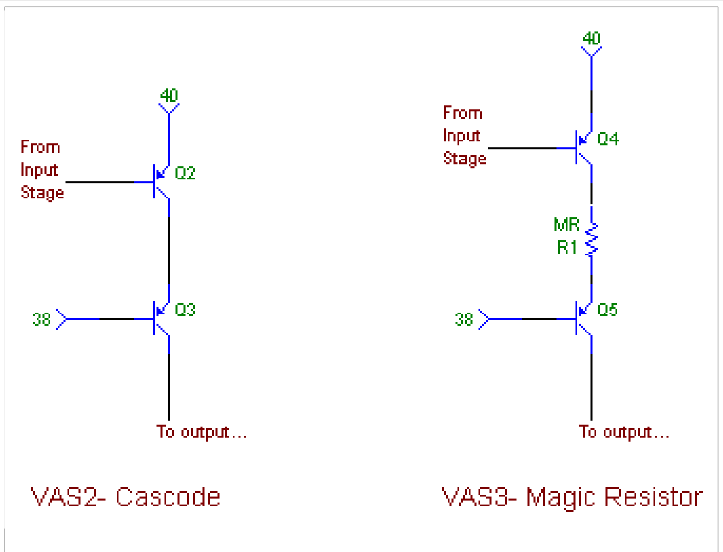

Fig. 5-3 : The Magic Resistor

We have a constant voltage **v** around the gain transistor (Q2/4), enforced by the cascode (Q3/5). However, the current **i** varies, even if by a small quantity. The Power **p** is not constant. Let **Vo** be the **v** without the Magic Resistor. then,

v = Vo - R*i

p = v i = (Vo - R i) i = Vo i - R i^2

Those who remember their algebra (?) will note that, if we choose **R = Vo/(2I)**, then at the DC bias point, little variations of **i** only affect **p** at the second order : power variations are therefore much smaller.

Tricky thing is, the emmiter resistor of the both transistors is part of R, so we can't use the above formula straight away ; a simulation will do nicely. VAS Transistor is BC560C, Cascode is BF470 :

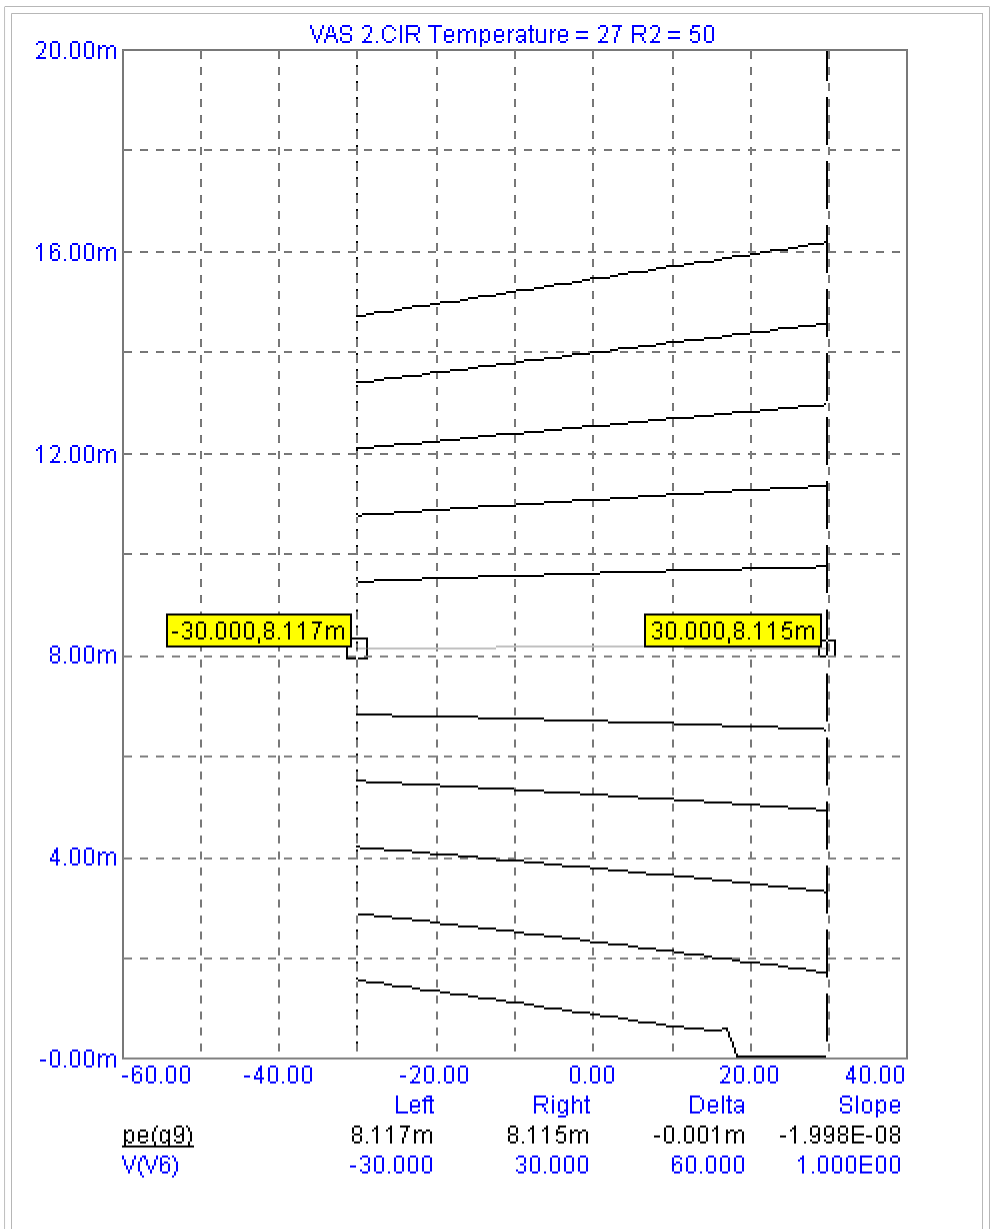

Fig. 5-4 : The Magic Resistor at work
Power dissipated in the gain transistor versus output voltage.

This set of curves is the power in the gain transistor versus amplifier output voltage, with the resistor stepped. 50 Ohms seems nice, as it gives an almost constant power. It works !

There should be no stability issues. Just in case, I put the resistor on the collector side, because the C-B paratitic cap is smaller than the E-B one.

Now, for a simulation of the complete amplifier. With the VAS as the only source of memory, the global memory goes from **4.56** to **0.4**. This is interesting, but it still swamps the memory of the input stage. Careful examination reveals the cause to be the thermal Vbe variations of the cascode transistor, which are enough to vary the voltage across the gain transistor. The Magic resistor can do nothing here.

What shall we do then... Well, obviously ;-)

Cascoding the Cascode

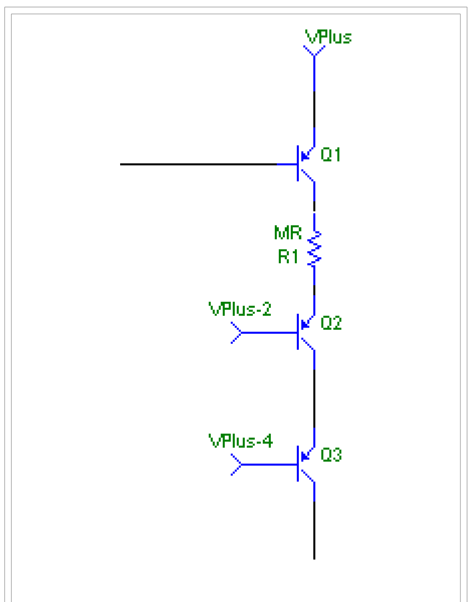

Fig. 5-5 : The Cascoded Cascode

Now, it is getting serious : VAS memory is down to **0.004**, and can be made to go as low as **0.002** with the magic resistor in place. The last part of the circuit that needs exercising is the current mirror :

Current Mirror	Memory
Perfect	0.01
Realistic	0.03

This can be taken care of easily by changing the emitter resistor in the current mirror for 1k instead of 100 Ohms. This gets us back to the value of a perfect current mirror.

Conclusion of the Theorical Part

We have taken a rather standard amplifier topology (Fig. 5-6) and transformed into a rather unusual one (Fig. 5-7), while reducing memory at least a 500-

fold. Now is the time for a listening test !

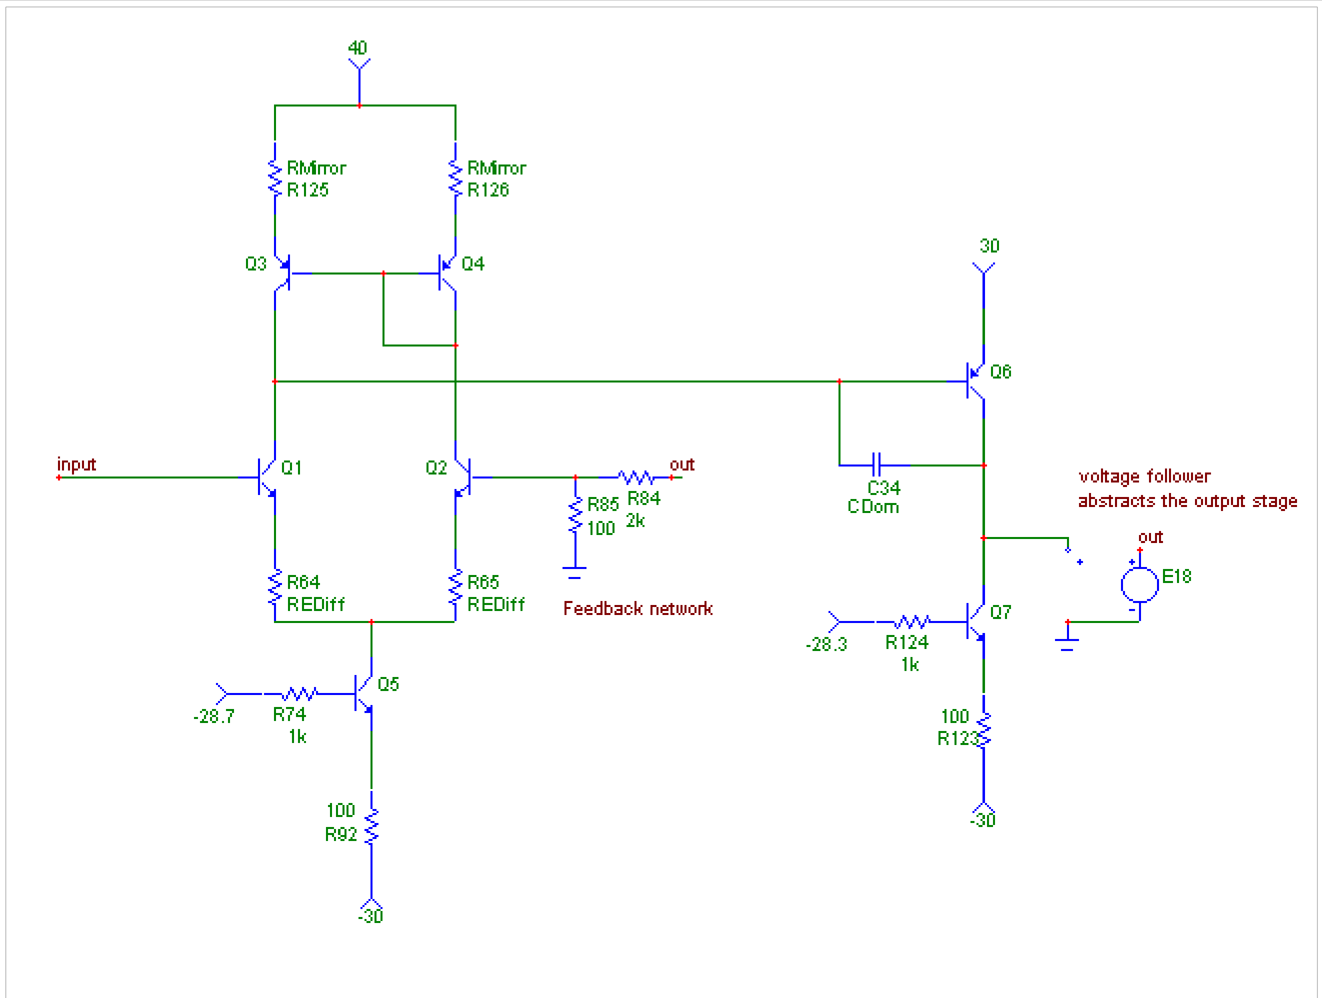

Fig. 5-6 : Before

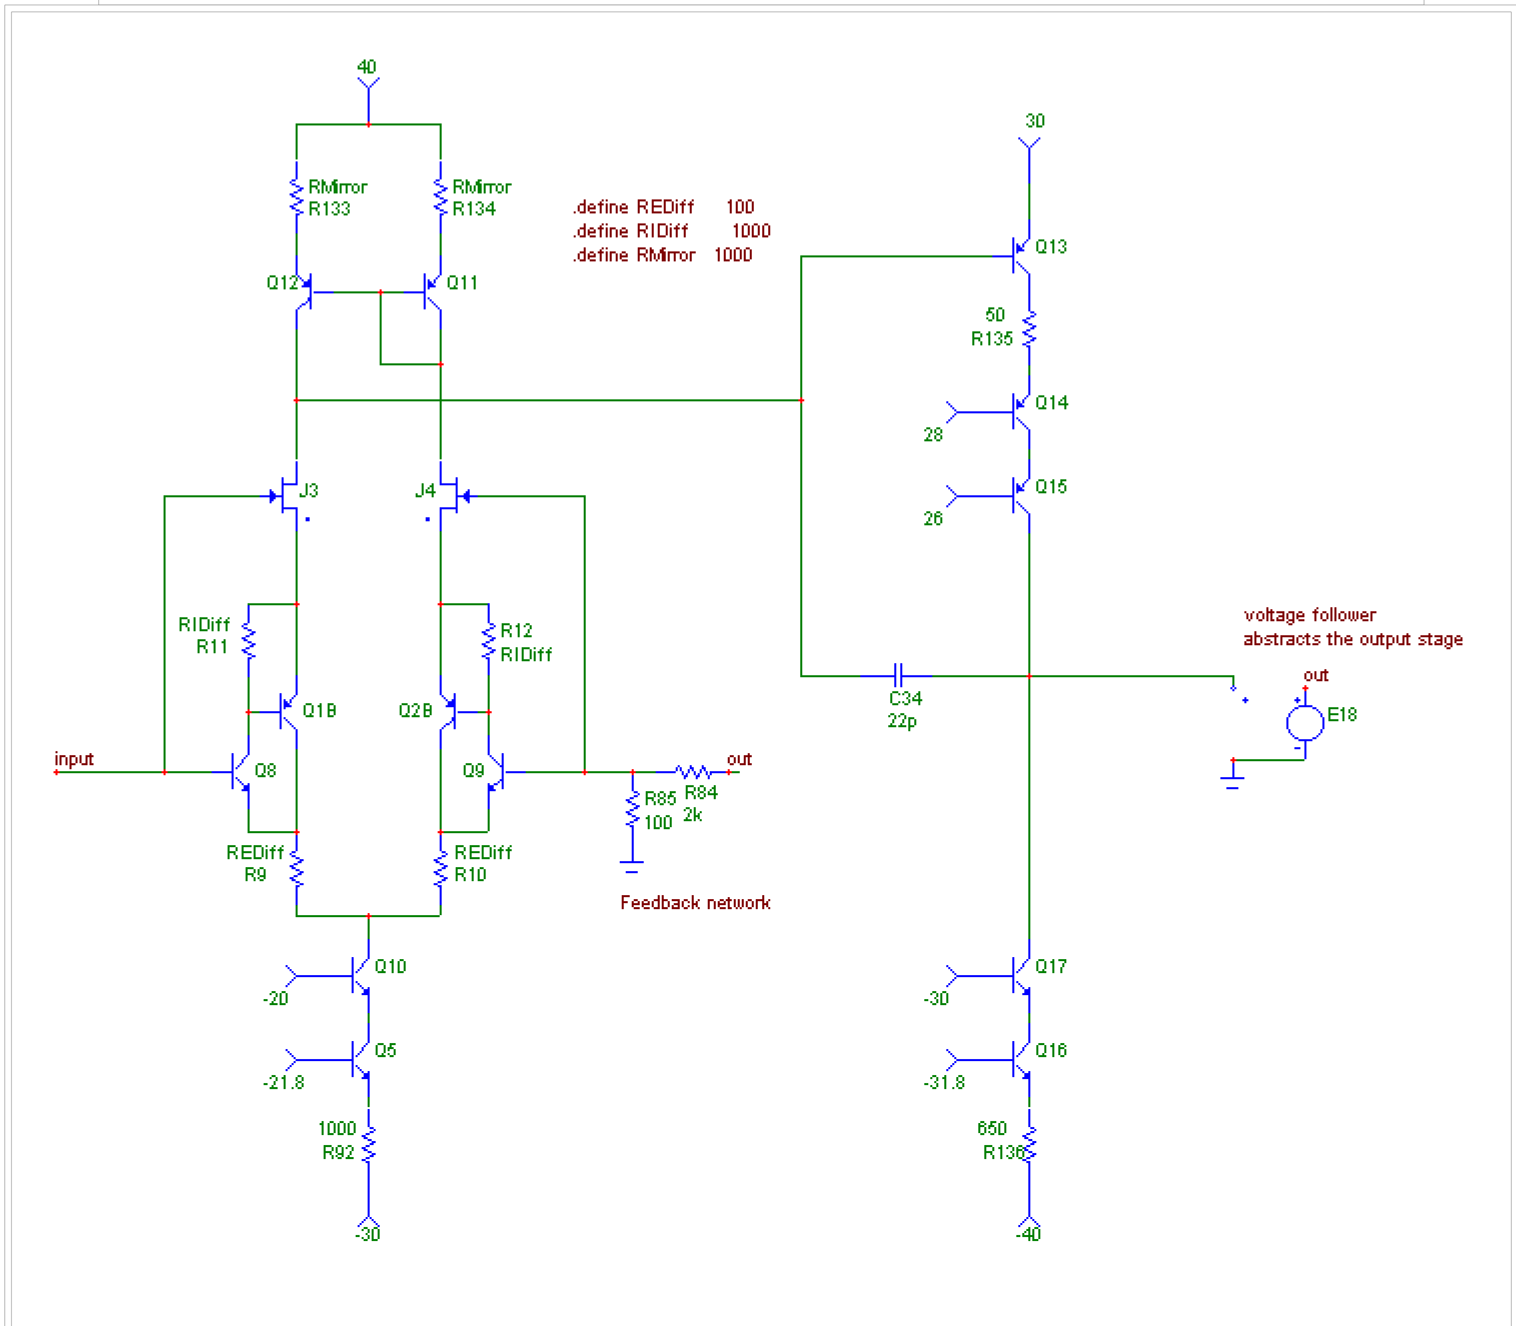

Fig. 5-7 : After

Memory Distortion - Part 6 : The Listening Test

It has amused me to read Douglas Self's *Audio Power Amplifier Design Handbook*. While this book contains a wealth of useful ideas, is really well written and scientifically-minded, it never mentions listening tests. What is the purpose of an amplifier, after all ? Do you listen to music to take the dust off the speakers, or to move your inner self ?

Self's book is compulsory reading for anyone wishing to design a power amp. I would ever have been able to design mine without this book. It is very competently written and interesting. However, we **will** proceed to the listening tests.

I used headphones for all my prototypes. There is no way I'm going to build a full power version of the amplifier straight away : when I burn 10 ouput transistors in two hours, they better be $0.50 ones that drive headphones, than $5 ones that drive speakers. Besides, the Sennheisers 580 allow me to hear many things that, on speakers, could have been masked by the imperfections in the room, speakers, whatever. And finally, my speakers are bi-amped, so that ended the debate.

So, the amp was coupled to a simple output stage powered by BD139/140, with drivers, a Vbe multiplier, and a buffer to protect the VAS from them. This is quite typical indeed. I tried single ended output stages too, but I didn't like the sound of them. Not enough balls.

Case	**Test Setup**	**Memory** (Simulated)	**What it sounded like...**
1.	Basic configuration (Fig. 6-1), with the current sources from the Advanced version (Fig. 6-2), and 1000R current mirror resistors.	Very high	Like solid state, harsh, grainy. You know the story. Bleh.
2.	Like 1, + Cascoded VAS	High	Better. Still solid-state, but less harsh. On the good way.
3.	Like 2, +JFETS in the input stage	Medium	**Much** better than 2, more fluid, more musical. Starts to get enjoyable.
4.	Like 2, with CFP in the input stage	Medium	Similar to 3, but maybe a little bit of a muffled, dark sound.
5.	Cascoded VAS + New input Stage	Low	Another big step forward. Music really gets enjoyable ! Solid-state hash is almost gone ! What remains could come from the CD player. The amp has almost no sound of its own, as the ambience changes completely from recording to recording. On some, it sounds genuinely Live. Goosebumps are there !

Scroll down please...

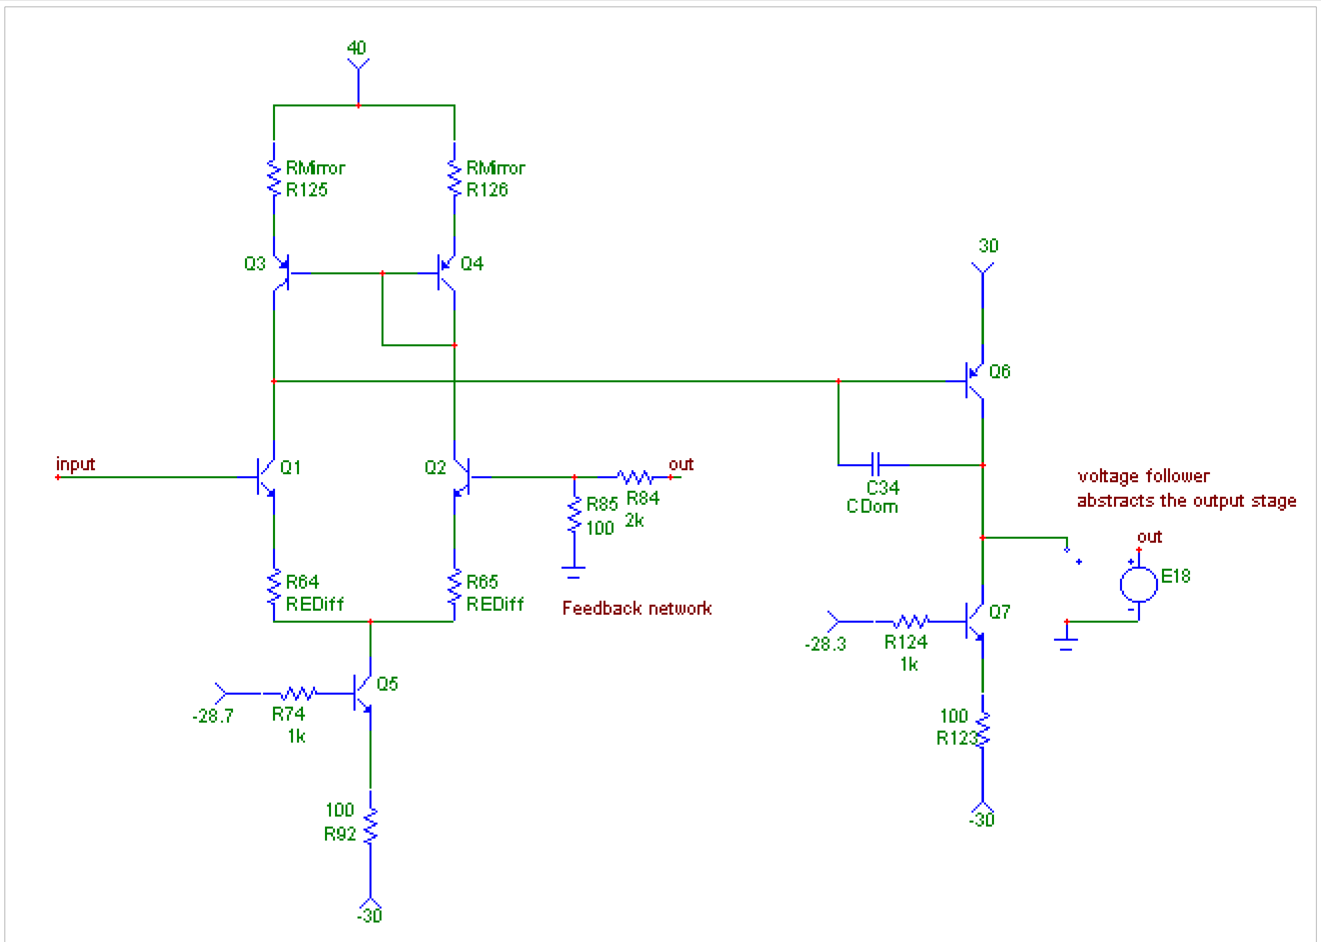

Fig. 6-1 : Basic configuration

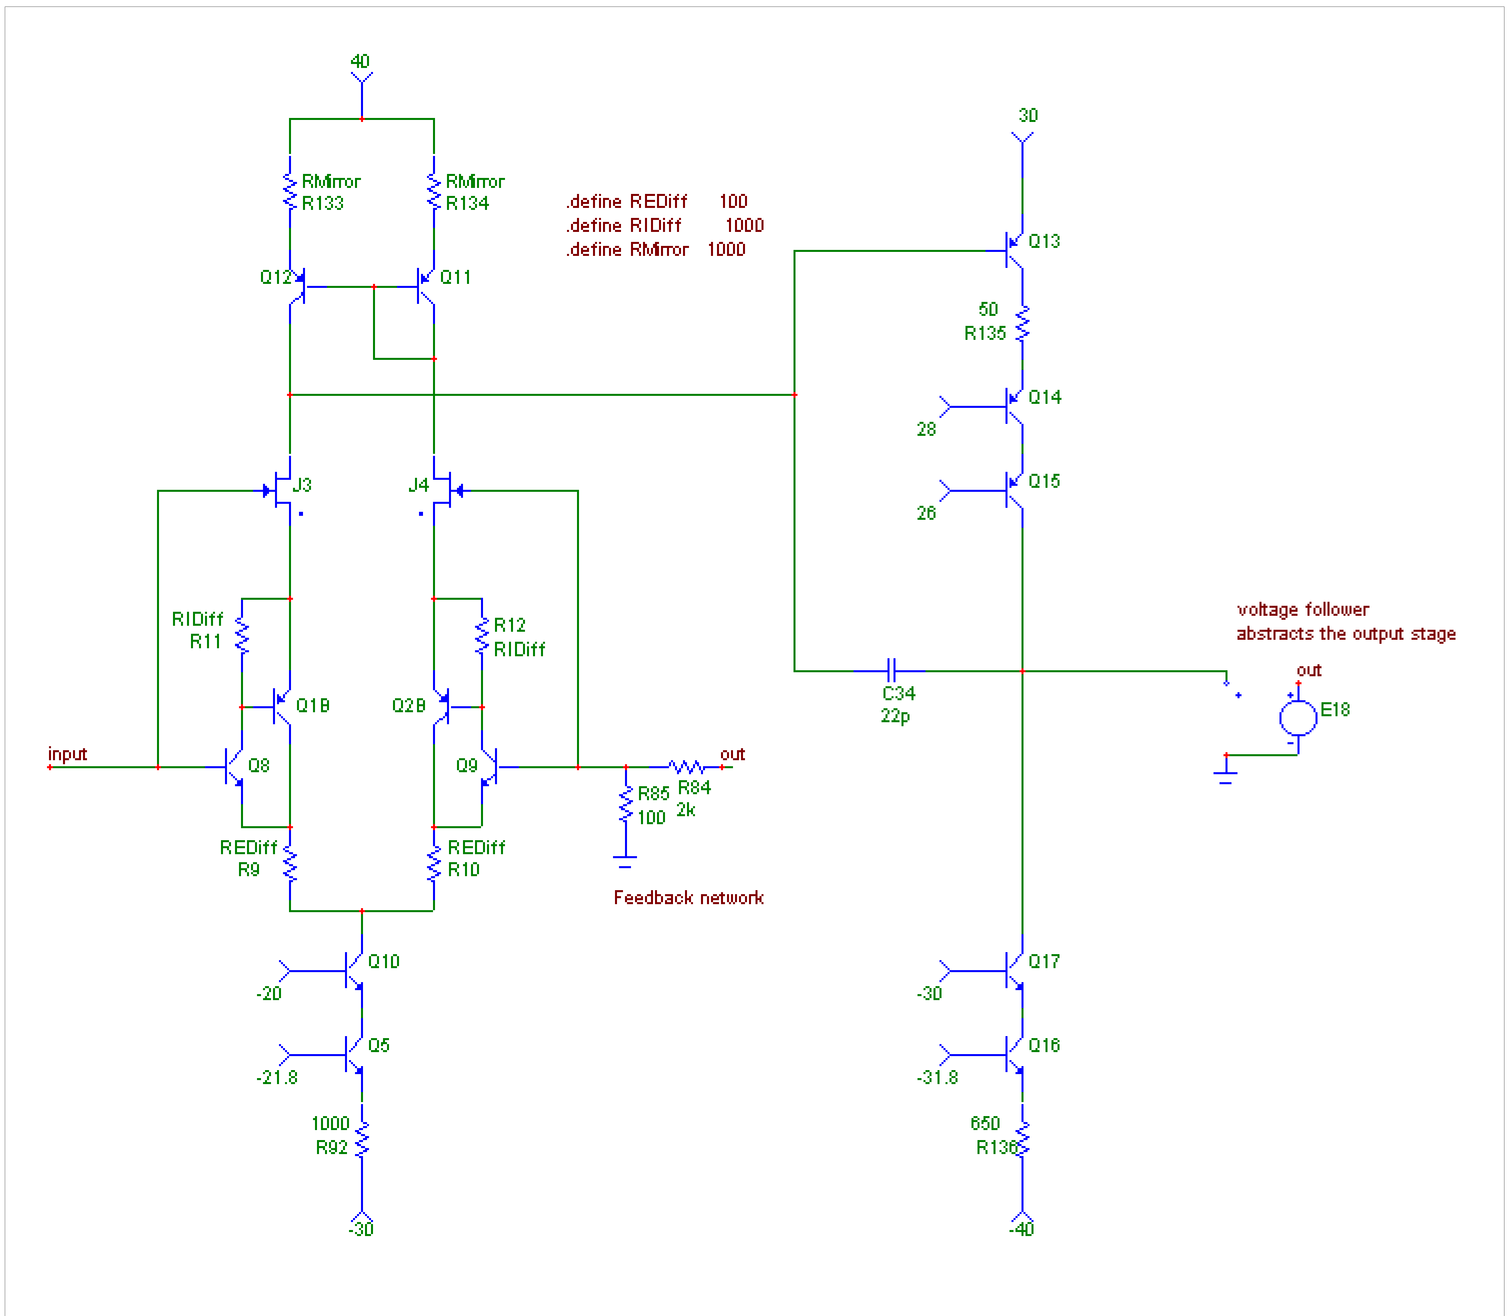

Fig. 6-2 : Advanced configuration

Temporary conclusion

So, the listening tests seem to confirm that minimizing memory is a good step towards improving subjective sound quality. Personnally, I have a problem here, as I can't listen to my stereo anymore. I just can't stand anymore how harsh, how solid-state it sounds. I spend hours listening to my prototype on the headphones instead, and I like it.

What remains to be done ?

Well, plenty. Two practical things, first :

- Testing the Cascoded Cascode VAS
- Adding a DC servo to get rid of the output DC drift (I use no input cap, so I can get upto 30mV of drift, which is bad with headphones).

And then, turning this prototype into a full-blown amplifier. It is going to take a lot of polishing.

The saturation problem is not solved yet. I will do another chapter on this, but as far as we are, now, when the amp clips, it gets as much memory as the very first prototype. This is very bad and has to be dealt with. I can't hear it with headphones, though, as I would be deaf before it clips... have you heard many headphone amps that can output 20V RMS ?

I am a little tired of all this writing. So, while I go rollerblading a little, please think about it and email me your feedback ! There are probably many typos, as I wanted this to be available early to keep the TNT amp project running...

Memory Distortion - Part 7 : Power Supplies

The effect of power supplies on sound quality has been known for a long time. In power amps, feedback is supposed to reject everything that happens on the power supply lines (the famous power supply rejection ratio).

The case of power amps is interesting, because most of them use unregulated power supplies. Therefore, when the amp needs to put out a lot of juice (on a transient, or when it clips), rail voltages will sag a little. How much will depend on the "stiffness" of the power supply (capacitor size), and wether or not the rectifying diodes are conducting when this current demand happens (in which case the current will come directly from the transformer).

Of course, moving rail voltages will not only affect dissipated power in devices, but also shift the operating parameters which depend on supply voltage. I tried to run some simulations with very pessimistic power supply sag (1 volt during 1 second), and I found that the memory effect *is* measurable, around **0.8** units in the scale we previously used, for the standard amplifier configuration. The New amplifier configuration that was examined in the last chapters shows no measurable memory.

So, here is yet another reason for regulating the power supplies in the signal stages. Dejan has been knowing that for what, 20 years ?.....

The Kick ass amp will, therefore, have two power supplies, one for the signal stages (fully regulated), and one for the output stage (unregulated, but with biiig caps).

Memory Distortion - Part 8 : More tests

Today, more listening tests. I wanted to try new VAS configurations :

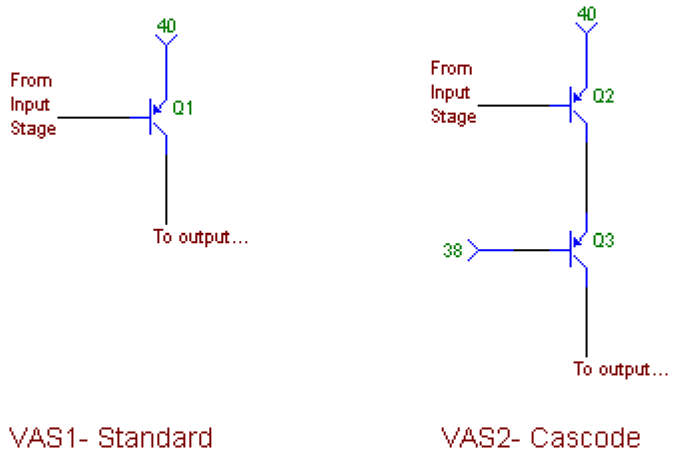

Fig. 8-1a : Simple VAS

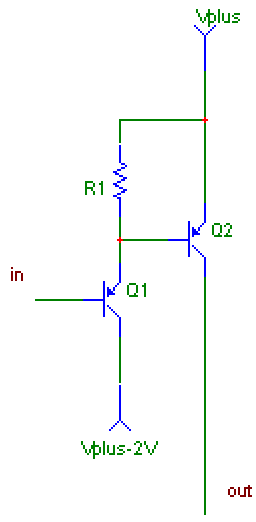

Fig. 8-1b : Beta-enhanced emitter follower VAS

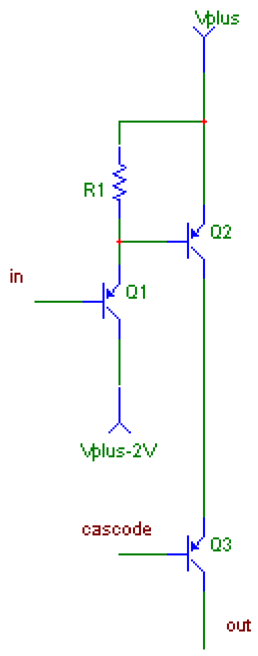

Fig. 8-1c : Same with cascode

The beta-enhanced VAS is interesting to compare to the Cascode because they have almost the same simulated memory values. However, the beta-enhanced VAS has more inherent memory, (because the power dissipated in the transistors is in no way kept constant). It *appears* to be the same as the cascode, only because it allows a much higher feedback factor, which in turn neatly hides all the gremlins.

On short term, A/B testing, the difference seems subtle, so I left the beta-e in for a few days to get used to its sound. A first it seemed to worsen the sound just a little, but after a while it seemed back to normal. After a few days I took it out and put the cascode instead. The change was *substantial*, more like getting another pair of brand new ears.

Now, when I switch, I realize that the beta-enhanced VAS, in spite of sounding good, does not stand a chance against the cascoded VAS of fig. 8-1a.

I went on searching how to get better sound from the beta-e, by trying to cascode it. This could give the best of both : a high gain and stable dissipation. I tried several configurations, which all did nothing except oscillate with total obstination. I gave up. If any one of you has ideas on this, please, email me.

Now, I tried to mess a little with the feedback factor. Have a look at fig. 8.2 below, and locate R142 (in parallel with the VAS compensation capacitor). This resistor makes a local feedback loop around the VAS and, by changing its value, we can transfer some part of the VAS gain from the global loop to local VAS linearization.

By the way, the schematic here is exactly what is right now in my high-tech laboratory (er, living room). It is my actual prototype.

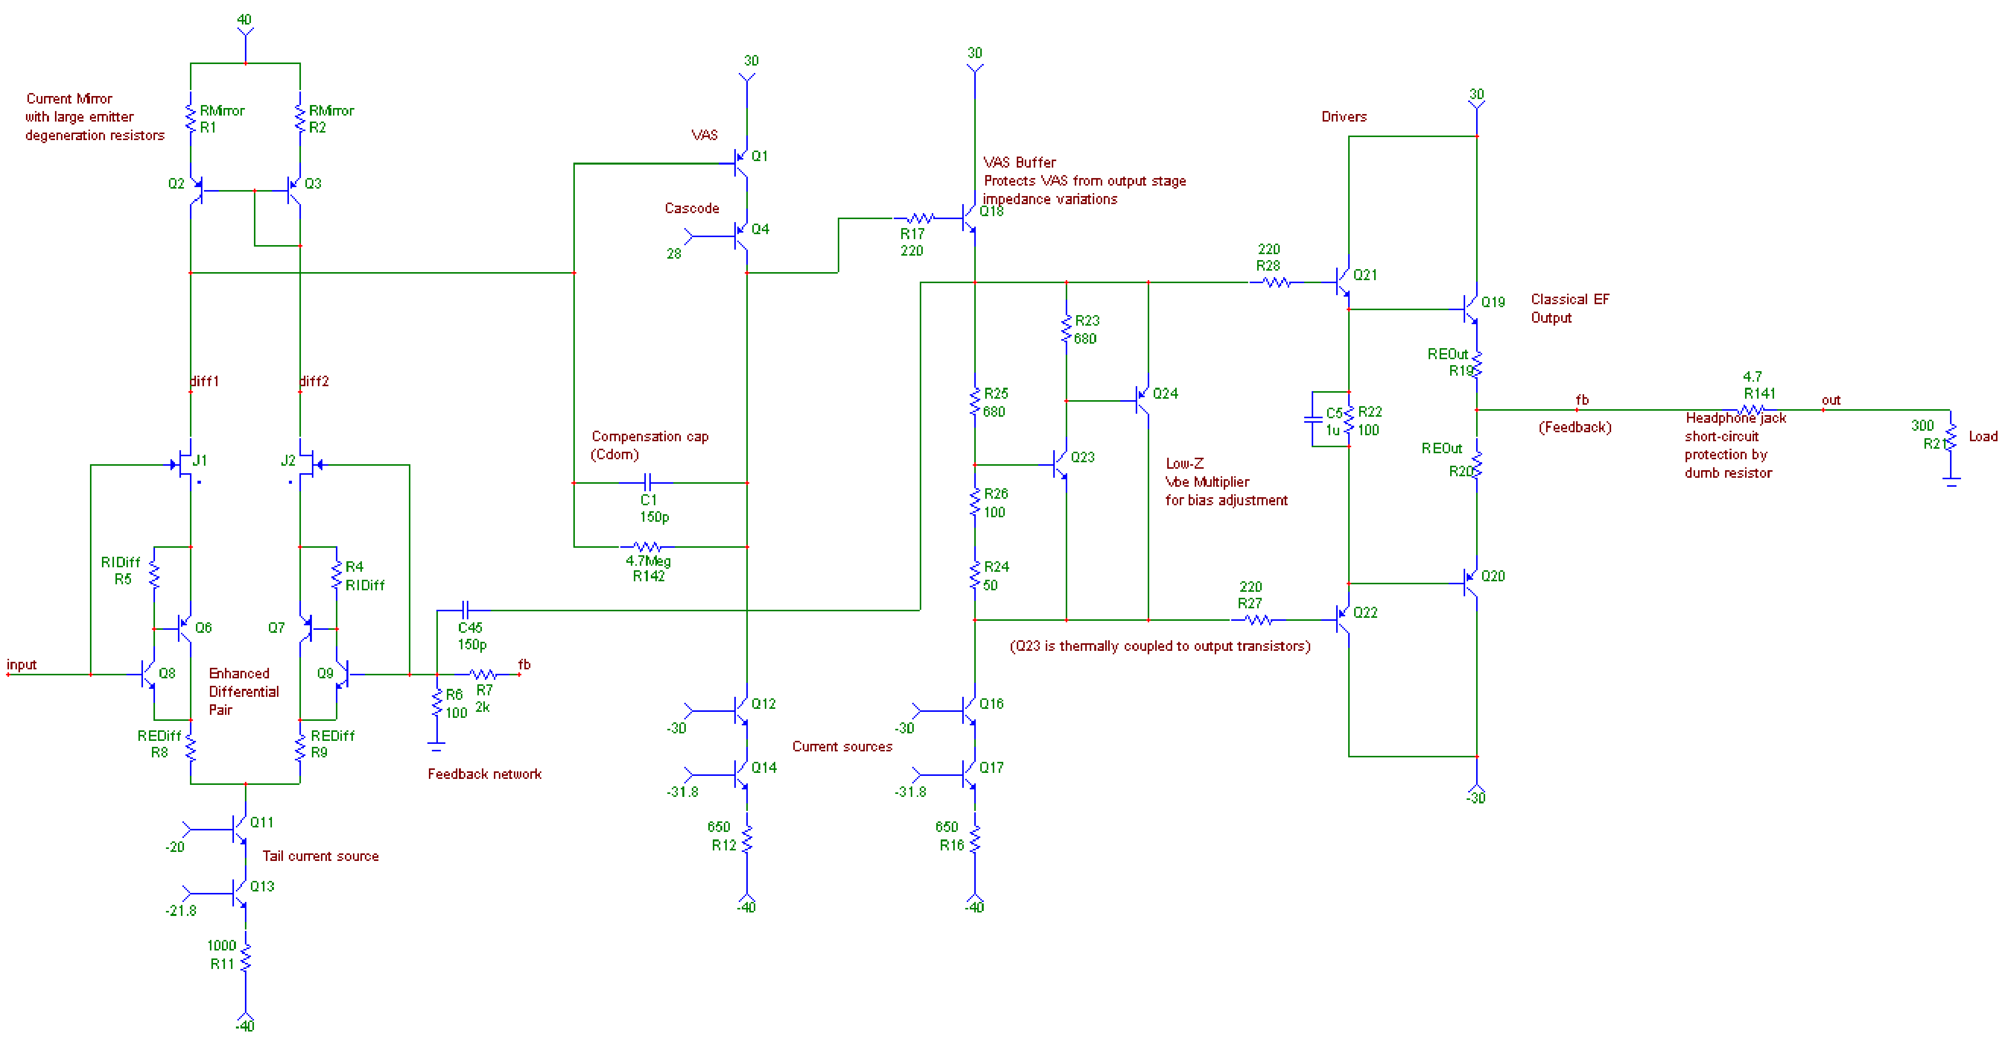

Fig. 8-2 : The Prototype

So, if we step R142 from 10K to open circuit, we get the curves in fig. 8-3. Horizontal axis in Hertz. Red is open-loop, black is closed-loop. Top is Amplitude (dB), bottom is Phase (deg).

The closed-loop amplitude curve gives us the closed loop gain, here 20 (26 dB). The space between open-loop and closed-loop gives the Feedback Factor. The values, at DC, are :

- R142 = 10K : 16 dB of feedback
- R142 = 4.7Meg : 67 dB of feedback
- R142 open circuit : 100dB of feedback (whoa !)

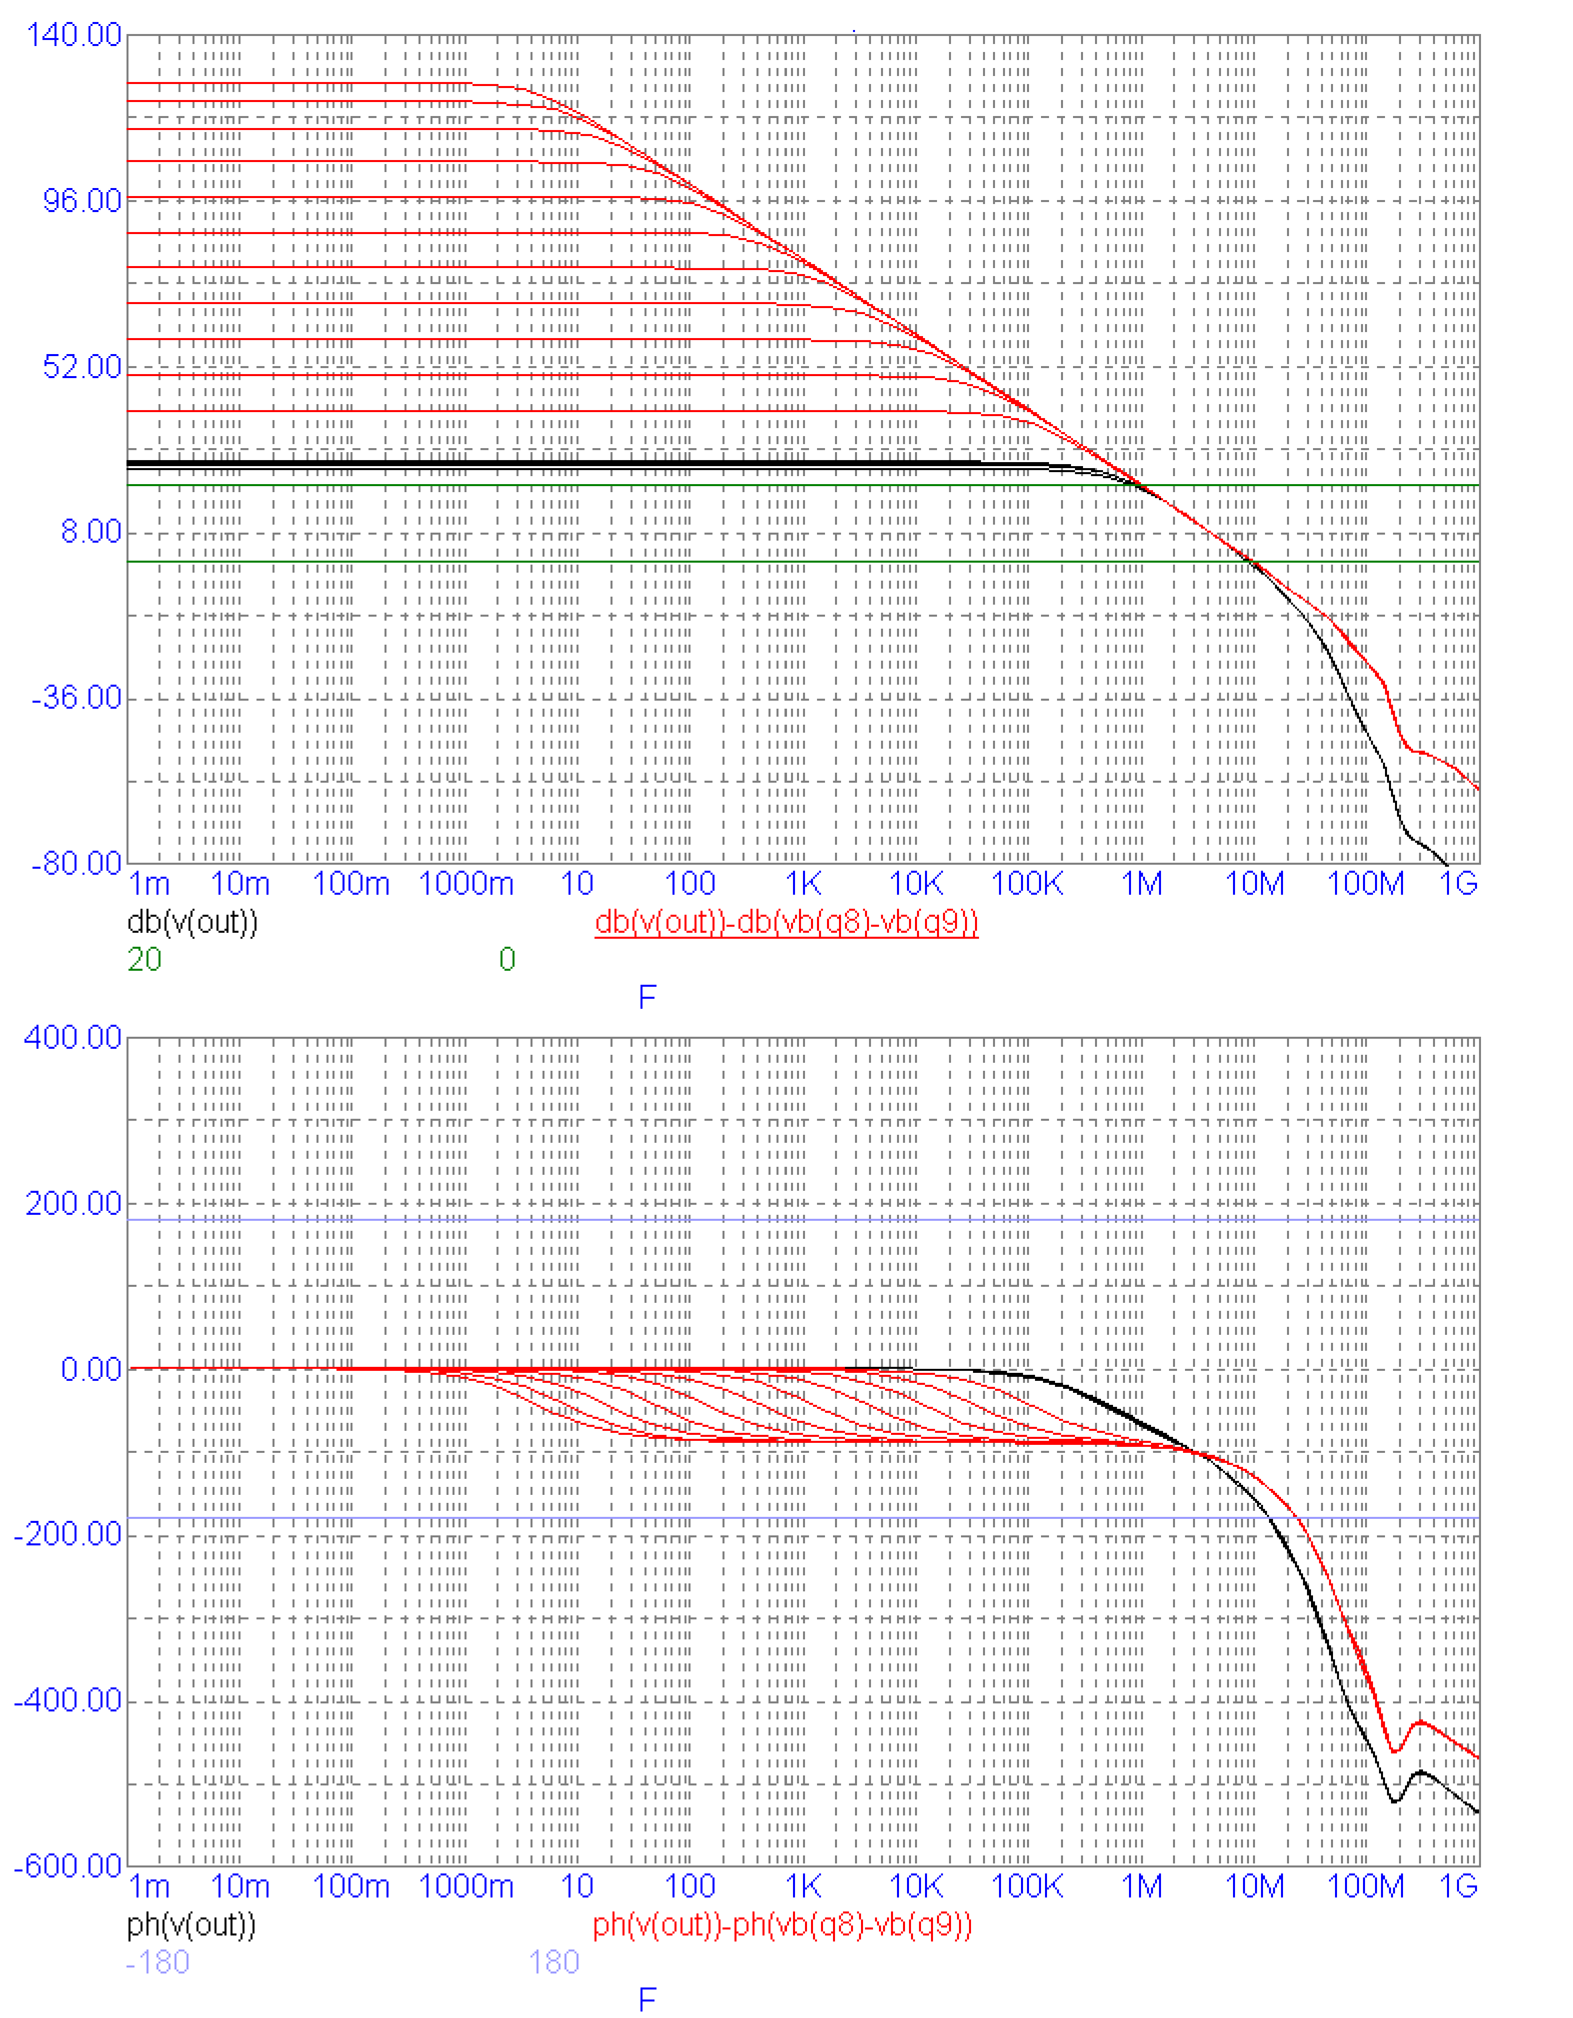

Fig. 8-3 : Bode plots for various feedback factors

I originally had 4.7Meg in, the exact reasons leading to this value being quite forgotten in my brain. So, in a provocative and bold move, I took out R140 . Now, the simulations might be a little "optimistic", as there are probably other factors happening in the real world, like un-simulated VAS collector loading for instance, which would lower the gain.

Now, how does this amp with a whooping 100dB of feedback, and a beautiful 3 Hertz open-loop bandwidth sound like ?

Whoa !

The ablation of this mere resistor made the amp get so much better, that I heard loads and loads and loads of stuff that I had never noticed, on my test CDs, the ones I know best, and that I have listened to non-stop while testing this amplifier in the past weeks. And it is not just "more detailed", it is more lifelike, more enjoyable. Its weak point is definitely the CD player driving it, as it reveals mercilessly the sibilance of the CD63SE, and the differences with the Harman DVD1, which does not sssibilate, but sounds a little dead sometimes.

For instance, the moment the singer is going to start can often be predicted by the minute changes in background noise, that happen when the sound engineer slides the mic volume pot up. On Alpha's "ComeFromHeaven", last track, the singer's mic signature is a slight hum with a particular tone. I don't talk about tracks where this is really too easy, like Diana Krall's "When I look into your eyes" album, which is really noisy ; no, on this album there is no silent background at all. In the same track, you can tell the samples from one another, and notice the drum loop is actually enhanced with a real guy hitting some drums from time to time, at a very low volume. This is also the first time I actually hear sounds coming from out of my head with headphones (not on all recordings).

The imaging and soundstaging is awesome. Everything is natural and pure. Everything changes completely from recording to recording. It can be **furiously violent** on Propellerheads, and reproduce chamber music in all its subtlety. It does not slaughter the voices of women. Bach's grand organ sounds like a grand organ (with the church reverb coming from all directions).

The really good thing is, if you want to count the tracks on the multitracks, you can, but then once you get used to these surprising things, it just clicks in you forget about it and listen to the music... it really flows, no aggressive tonalities jump at you. It is Musical. The details are there in a natural way, not jumping at you, just being there instead of being erased by the electronics... Let me do the world's shortest review :

It kicks ass !

Next things to do :

- How can I get more feedback ? More than 100dB ? Hummmm...
- How can I get this beta-enhanced cascoded VAS **not** to oscillate ?
- What happens on Clipping. I think I will make the amp clip and then use an attenuator to hear just the effect of clipping, but without getting deaf. Probably the only viable solution will be to pre-clip the signal before it enters the amp, with diodes or whatever, so that the amp itself will never clip. The ideal would be to do it in the digital domain, of course.

Problem is, I'm short on ideas now. What shall I do next ? Well, I still have 400 albums to re-listen to...

- Comments from readers.

Pierre,

Excellent, excellent article. I only read it lightly, but found nothing to argue with. I too believe one of the "differences" between SS and VT amps to be just how they are affected by instantaneous temperature changes in the active devices.

These are ideas that I have intuitively considered to be at play, but never bothered to attempt to measure.

I Look forward to more.

Marvin Match

University of Utah

Thanks for the comments Marvin ! This is what I hope for : wake up people who have been thinking about this kind of things before, and think together about it. Read the new articles, we'll talk about it.

Pierre

I find some of your comments especially interesting- particularly that the brain compensates for or ignores the constant distortion of the SET. The function of the brain in what we think we hear is to me the big black hole wherein probably lies a lot of the apparent difference in peoples perception of sound.

I am somewhat surprised that you write off TIM as I would think that what you describe is a very valid TIM mechanism. ie change of operating points with the possibility of increasing distortion. Particularly if the amplifier is direct coupled to the output with resulting change in output bias. The distortion of the output stage is very dependent on bias settings. regards pat

Some research reports that the ear has lots of distortion by itself. I might be speculating, but this might explain why the high distortion figures of SET amps don't seem to matter. Just like if the brain could filter out large amounts of low-order distortion, on the condition that the spectrum be stable, which is not the case with transistors.

You are also positively right about TIM ! I think that conventionnally TIM is associated with slew rate limitation, but this memory effect could also be labeled transient intermodulation as it is caused by the signal itself... It is a matter of vocabulary actually. Personnally I'd rather call it memory to differentiate from slew rate effects.

Thanks,

Pierre

I have now read your ideas on the SS amplifier on the page http://peufeu.free.fr/audio/ and found that your results correspond very closely to my efforts.

The use of a cascode stage to drive the output is very important as well as the use of good CCS for both the input stage and VAS.

I will try the magic resistor tonight and see what I can hear. You comment that during listening tests you could still hear the hiss and suggested it may come from the input signal, I would hazard a guess and say it is coming from the zener in your CCS.

Thanks ! Working alone, and not having other listeners at hand to confirm what I thought made me a little nervous. Was I hearing things ? Now, I receive emails that confirm from experience, that structures with lower memory have better sound....

I would suggest that you look at using the Philips BF469/470 as they are much better sounding and do not inhibit the slew rate as do the BD139/140. There is also the MJE 340/350 series but they also inhibit slew rate but have better driving abilities than the BD139/140

For small signal I use BC556 or BF421/423 which seem to be similar apart from the VCEO.

Power transistors : I have tried the Toshiba 2SA1302 /2SC3281 but found them to be very unstable over about 35 Volt rails. I am now using the MotorolaMJL21193/21194's with good results.

MOSFETs do work well but are not as good sounding as the Bipolars and I have not worked out why yet.

Maybe this is because MOSFETS are much less linear than Bipolars... their crossover region looks very bad unless biased into class A. Even then, they are less linear than bipolar.

I would also stay away from the use of CFP stages as they tend to get unstable and the local feedback errors generated cannot be controlled by the global feedback.

I tried CFPs too, they seem to like oscillating very much. Stabilizing them (with caps) does not seem to do good to the sound. The good old EF stage is IMHO the best compromise.

Another worry is what effect the normal VBE multiplier has in the bias arrangement for the output stage, as I think this also has a detrimental effect on the sound due to it trying to present a stable voltage during transient conditions as well as providing bias stability during thermal changes to the output stage. The use of a capacitor helps but then you get into capacitor problems again. I would prefer to use a simple resistor with the thermal characteristics being taken care of in another way. Russell.

To be honest I don't know if the memory occuring in the output stage is important. I have got to make other simulations on this. It is actually an other form of memory than the one I already studied, as it is not the DC drifts that are annoying (these are verry little), but rather the crossover distortion being modulated by the quiescent variations caused by transistor heating/cooling. The cure might simply be to use large emitter resistors.

I think that the Vbe multiplier is slow enough... The transient is finished for long once its thermal effects start to arrive at the sensor transistor. The memory effects I simulated are much faster (ie. at the second scale, similar to the beat of the music, or even faster, ie. the bass notes below 100 Hz). Output stage memory and interaction with the Vbe multiplier is at much slower speeds, like the alternance of loud and soft moments in music. This might be very important, too, and the easy solution appears to be Class-A, with all assorted heat problems.

Unlike memory in the rest of the amplifier, the output stage has been quite studied before. I will try to simulate it nonetheless...

Thanks for your feedback !

Pierre

These are messages I got on the Audio Asylum, and their replies.

Thanks for sharing your facinating work. I do have some questions -

1.Your Vas circut ,cascode or cascoded cascode, would, I believe, have a very high output impedence. With the addition of an emmitter/source follower you have a defacto, functional op amp. A wonderful building block. Have you investigated the succeptability of followers to memory distortion, and its cure? If so could you update us on your findings?

2. Have you expolered the application of your findings to balanced/differential circutry? Hope your original work bears fruit and look forward to the day I can buy a 'minimal memory' op amp from your multinational conglomerate.

I failed to take note of the cap providing local feedback around the (cascaded-) cascode. Could you comment on the output impedence and load driving capabilities (modeled or actual) with the circut as it now stands?, as well as the previous question on followers?

The VAS :

- output impedance : intentionally very high (I want to squeeze out the most gain from it, so the impedance has to be highest). Impedance decreases gradually with increasing frequency, as feedback through the Cdom cap comes into effect. This is the standard method of VAS linearization through Cdom.
- drive capabilities : almost none at LF
- hence the buffer transistor !

The memory in the **buffer transistor** is quite negligible compared to the VAS transistor. Imagine the VAS is confronted to a 50k load (base resistance of the buffer), with a 10mA current sink. Heating the VAS transistor by one degree C will vary its Beta by 3%, which means a drift of 10mA * 3% * 50k = 15 volts that will have to be corrected by feedback. Heating the buffer by 1°C will only make it drift 2mV... The essence of the thing is that memory is multiplied by the gain of the stages downstream ; the VAS has a lot of gain, the follower has none.

If you use an emitter follower for amplification (with for instance 10 Ohm from emitter to ground and 1k from collector to V+) as is done in an input diff pair, you're into trouble, because the Vbe variations are amplified by the gain of the stage. They have to be compared to the input signal, not the output signal.

With a buffer and a small output stage, it is actually an opamp, even if a complicated one... It could be used for anything actually.

Early Effect :

I don't think memory distortion is related to the early effect. This one does not depend on temperature (or so I think). This is more about Vbe variations...

Headphones :

I have Sennheisers 580... are there any better headphones besides the 600s ? Electrostats do not apply here, thay usually have their own tube amps.

Aim at the sore spot

I was trying to figure out which transistors of the whole amp had the most influence on the total memory distortion. Then, the answer struck me as obvious : let's calculate the following two values :

The Annoyance Factor

- Considering an open-loop amp to which we apply just the right DC input voltage so as to set the output to 0 volts. This is easy to do in simulation, much less so in real life. All results here are in simulation.

  In open-loop, the amp is very sensitive, so we can measure **F = dV(Output Open Loop)/d(Temp)**, the variation of open-loop output voltage caused by a 1° variation in the tempearture of a specific transistor. **F** is in Volts/°C.
- Then, in closed-loop, we can measure **H = d(Dissipated Power)/d(VOut Closed Loop)**, which is the variation of power dissipated in a specific transistor when the output voltage changes by Volt. **H** is in Watts/Volts.

Therefore, for each transistor, we know

- **H** : which tells us how much this transistor will heat up (as Temperature variations are proportional to dissipated power variations) ;
- **F** : which tells us the gravity of the drift induced by this transistor heating.

The product of these two values : **A = H*F** therefore gives an idea of how Annoying this transistor will be relative to the thermal drifts of the whole amplifier.

Example 1 : classic Lin amp

click to display schematic (hint : use shift-click)

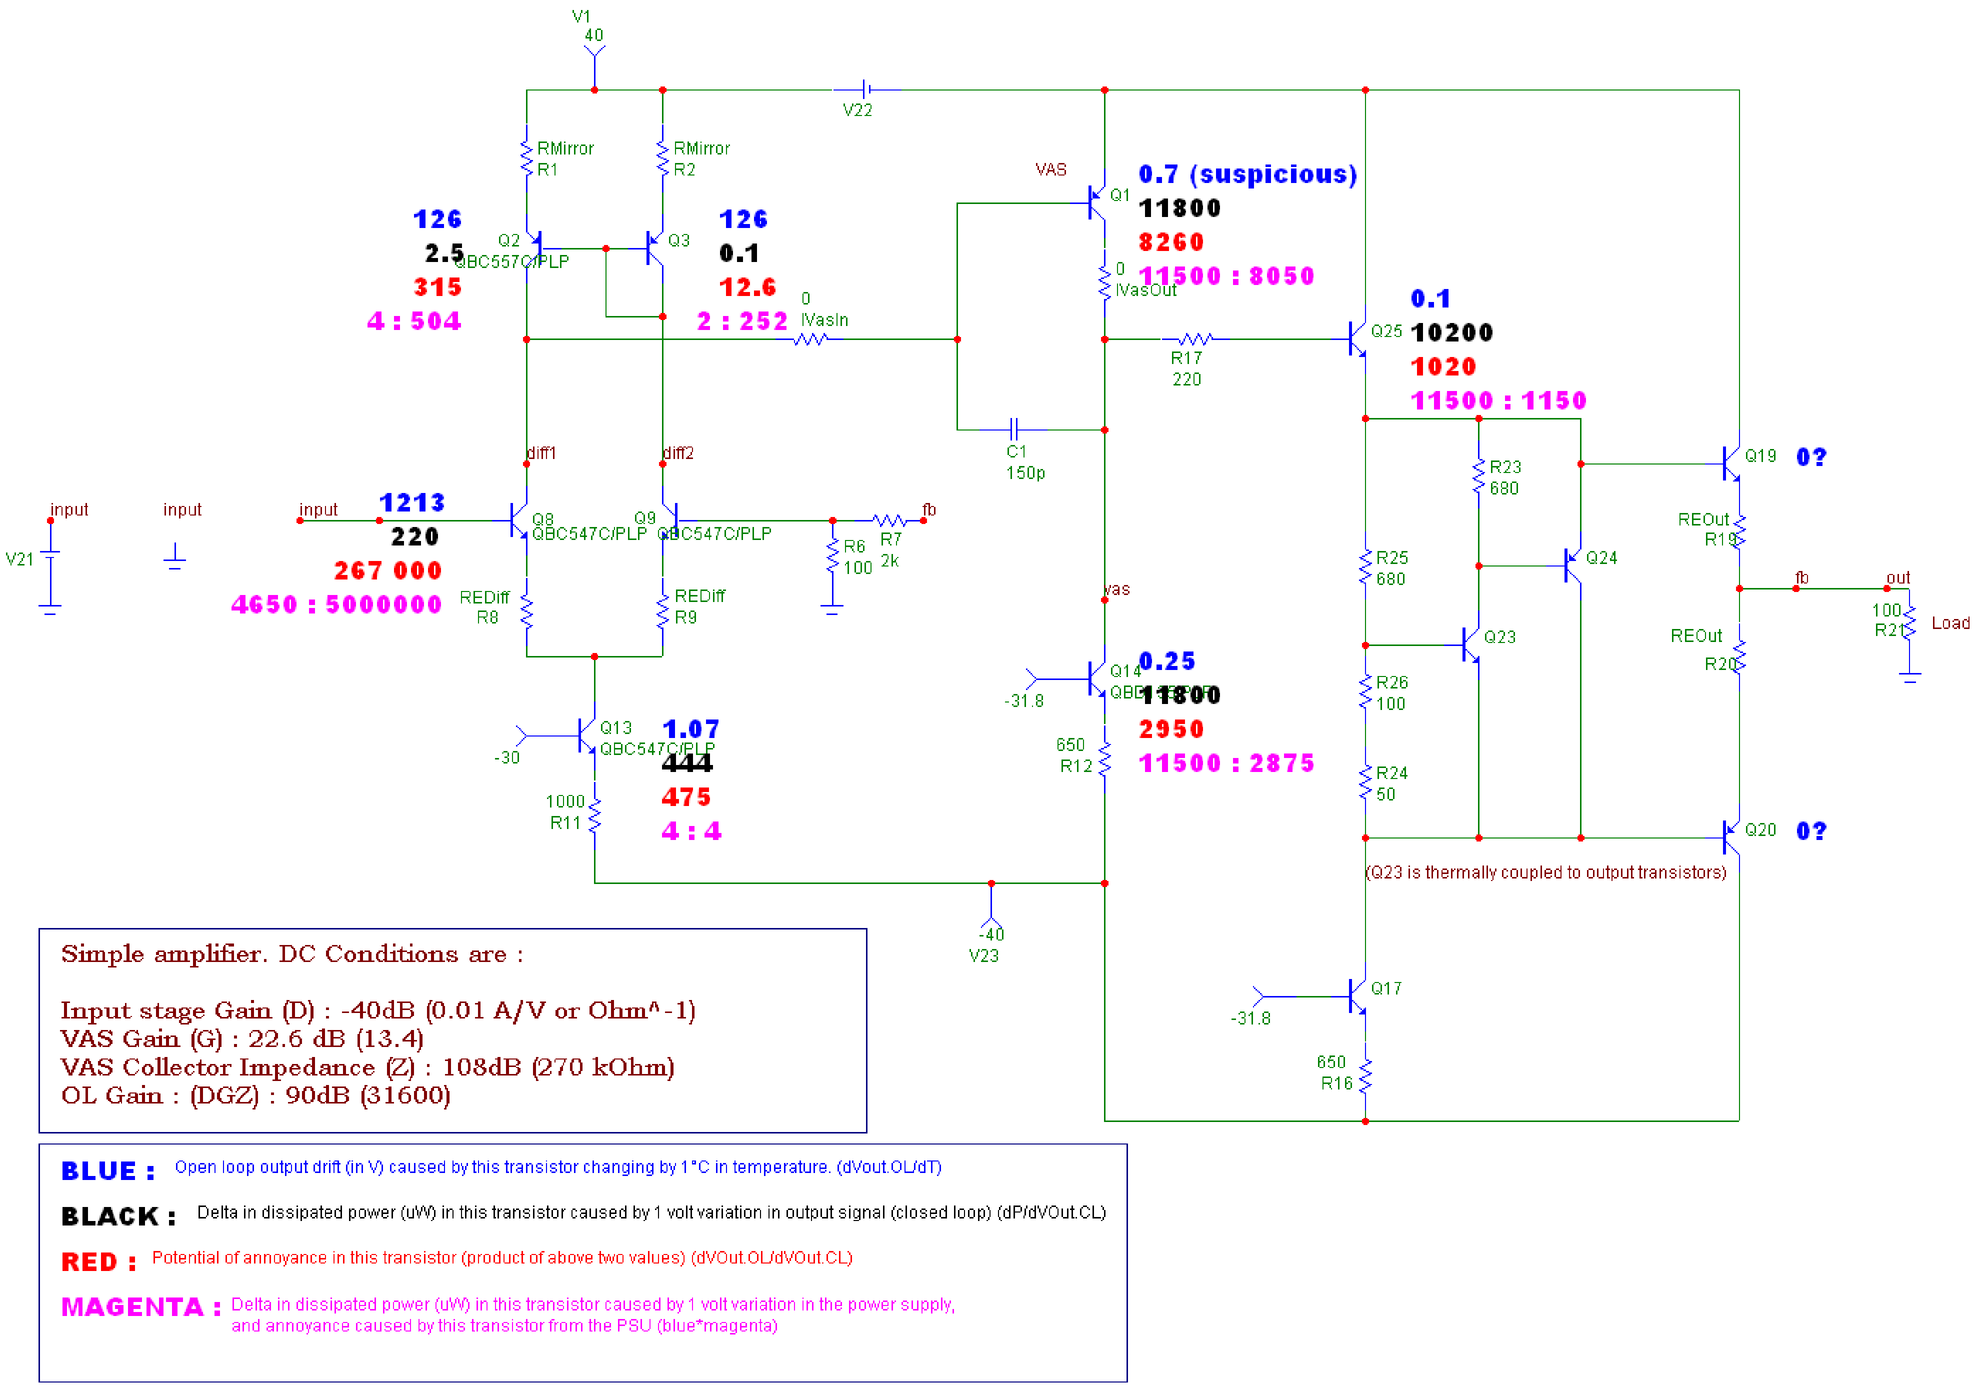

Sorry for the huge graphics (I always work at 1600x1200). If you use Opera, hit "-" to zoom out. If you use an imperialist browser made by a monopolistic cartel, or a product of a defunct company, hit ctrl-alt-del.

Well well well, this is a standard sand amp, 90 dB open loop gain, 24 dB gain, 66 dB feedback at DC. Close to each transistor, I wrote the corresponding H, F, and A.

- **F**, in blue : open loop output drift in volts if this transistor heats by 1°C.
- **H**, in black : power delta for 1 volt Vout change, in micro-watts
- **A**, in red : annoyance factor, ie. product of above two values.

Now it is pretty obvious that the most annoying transistor is in the input pair, with some help from the VAS and its current sink.

The numbers for H are pretty logical, just think that Power=Voltage*Current and try to explain them.

The numbers for F get worse the closer you get to the input, because they are magnified by the gain of the rest of the amp. ie. a drift occuring in Q25 will not be multiplied by anything, whereas a drift occuring in the input pair will be multiplied by the full OL gain.

In Magenta, I wrote the results for the same technique applied to the power supply. Let's call **H'** the power variation caused by a 1 volt variation in supply voltage : d(Power)/d(Supply). Then, **A'=H*F'** relates to the sensitivity of the concerned transistor versus the power supply. The two numbers in magenta are H' : A'.

We can see the input stage is extremely sensitive to power supply variations.

As a side note, the thermal effects in both transistors of the input pair cancel each other, but not completely. The sore spot is therefore not only in the input stage, but also in the VAS and the current sources.

Example 2 : thermally compensated amp

click to display schematic (hint : use shift-click)

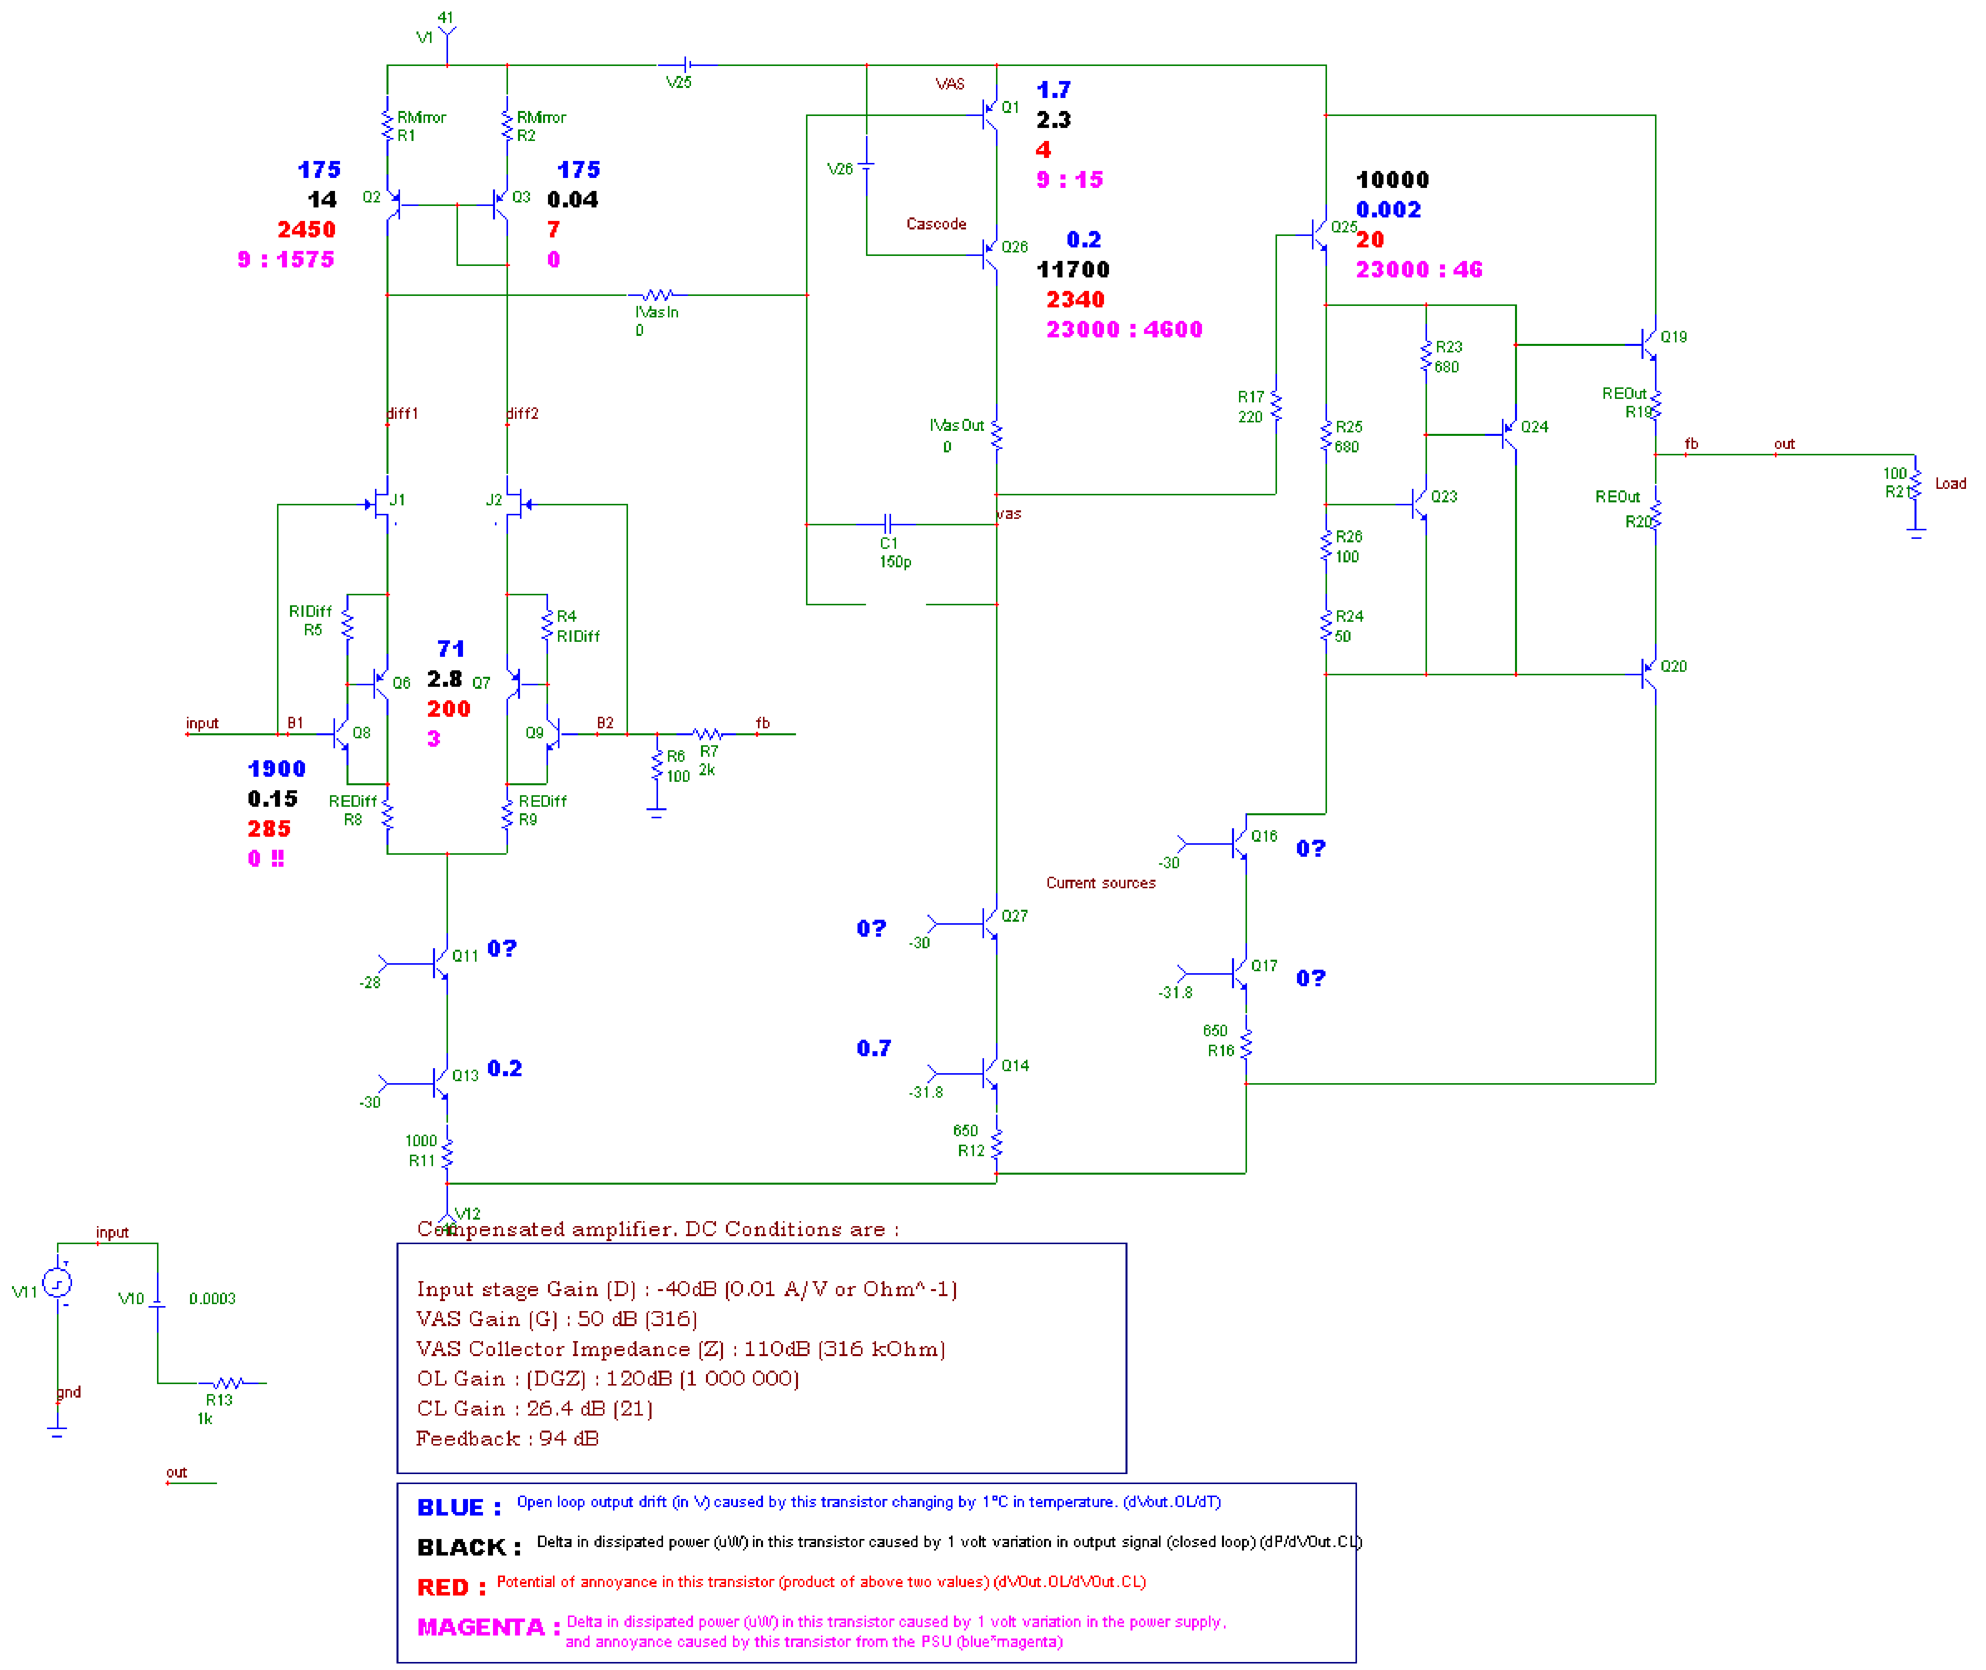

The steps taken to operate critical transistors at constant power make the input stage almost imune to drifts : its annoyance factor dropped dramatically, both relative to signal (red) and to the power supply (magenta).

The current sources also benefitted a lot from cascoding. Same for the VAS.

We can see where the new sore spot is, and where to direct the efforts :

- The current mirror (I'll add a simple JFET cascode, à la Borbely)
- The VAS : under investigation.

A major interest of this method is that the buffer transistor is shown to be quite innocent, where I would have thought the contrary.

Maybe this thing is going to get finished one day ?

All in all, the compensated amp sounds good (euphemish). 94dB feedback, why not ? I have no preconceptions on that.