VCAS INVESTIGATED PART ONE

Ben Duncan outlines where VCAs belong in the scheme of things, discusses their origins, and the circuit topologies adopted by competing VCA makers

While VCAs appear with increasing frequency inside pro-audio equipment, readable literature on the pros and cons of competing VCA techniques is so scarce and widely scattered, that few sound engineers (or equipment designers for that matter) have had time to thoroughly 'get their head around' the different approaches. Yet without some degree of technical insight *and* overview, users (let alone spectators!) are almost bound to develop irrational prejudices.

VCAs (Voltage Controlled Amplifiers) provide the electronic control gain and attenuation that's at the heart of most audio processors and all kinds of automation systems, remote-controlled faders, panners and equalisers. Still, casting a glance at **Fig 1**, there's more than one way to implement electronically controlled amplification and attenuation.

The family

Looking at the right-hand branches of the evolutionary tree in **Fig 1**, the incentive to develop audio VCAs arose out of the limitations of VCRs, which are much older and simpler (in concept, if not in practice). Both kinds are related through ohms law: VCR is just a voltage controlled resistance.

On its own, a VCR provides attenuation alone, but placed in a feedback loop, it can just as easily control the gain of an amplifier. Implementing the variable resistance implies using one or other kind of FET, or an LDR. Over the dynamic range required of a fader, and assuming a potentiometric arrangement, both kinds of element have awkward, uncivilised non-linearities of the sort which can't be wholly overcome—even granted fairly convoluted support circuitry. Ever since workable VCA topologies were established (in the mid '70s) VCRs have been left to applications which aren't fussy about limited linearity and a restricted operating range, ie limiters and budget FX processors.

The branch below is old, but a smattering of new growth affirms it's still living. On it are motorised pots. Pots (potentiometers) and motors are both Victorian inventions. Hooking them up to late 20th century logical electronic control works, but to Marlowe, it doesn't seem so elegant. In common with other artefacts of 19th century engineering, the combination is characterised by simplicity, no great regard for size, energy consumption, or mechanical noise. To most readers, the long-term disadvantages of pots and stepper motors won't need much spelling out. Computer disk drives are hardly a good advert. Less well known is that (supposedly) SOTA pot manufacturers are still unable to manufacture ganged stereo faders (whether rotary or linear) which maintain channel balance within ±½ dB over the 30 dB+ span of everyday monitoring SPLs. When it comes to souping-up mechanical components, even laser trim has its limits.

The tree is lopsided: its left side has just one main branch, which is young and bears many buds. The contents of this branch are fundamentally different. While the devices on the right hand branches exhibit essentially infinite resolution ('you can put it where you really want it'), the left branch is all about producing gain and attenuation changes in discrete steps. Even for relatively slow processes, like automated faders, glitch-free gain changes demands near equal steps of below 1 dB. For a fader covering 70 dB, it calls for (at least) a 12 bit multiplying DAC. Dedicated digital attenuator ICs have existed for over a decade, but their performance is still regarded with some suspicion. In the realms of customised pro-audio and money-no-object domestic hi-fi, volume controls have been built from discrete FETs. Granted a great deal of logical ingenuity[1], the number of FETs needed to switch through 70 dB in ½ dB increments is less than you'd imagine.

For remote EQ, the ear is less acute. In the world's first automatic tonal compensator[2], dynamic increments as big as 5 dB aren't readily audible at certain audio frequencies. Twenty years after their inception, FET switches which don't go *blahht* (particularly when the wick's turned up), don't distort at high levels, don't raise the noise floor unacceptably and don't require convoluted circuitry to support the cancellation of resistance modulation, are still scarce—if they exist at all. Finally, in processors, where the rate of gain change required is commonly rated in dBs per mS, the requirement for glitch-free stepping is even more stringent. Did someone say zipper noise?

Analogue Computing

Audio VCAs were originally contrived in the late '30s by Western Electric for the compressors and limiters needed for film dialogue recording and broadcast transmitters. Further developments occurred in the '40s and '50s, when VCAs were used for analogue computers, at a time when digital machines weren't the fastest way to solve multi-order differential equations for the trajectories of Inter-Continental Ballistic Missiles, or the problems of fluid mechanics posed by nuclear reactors and chemical plants.

Since then, analogue computers have grown and grown, except nobody calls them that any longer. They come under the heading 'Analogue Functions' or 'Instrumentation'. In the language of boffins who want to monitor industrial processes, VCAs are multipliers or dividers. They are analogue's way of carrying out computation. Couple one VCA to another, and you've built a ▷

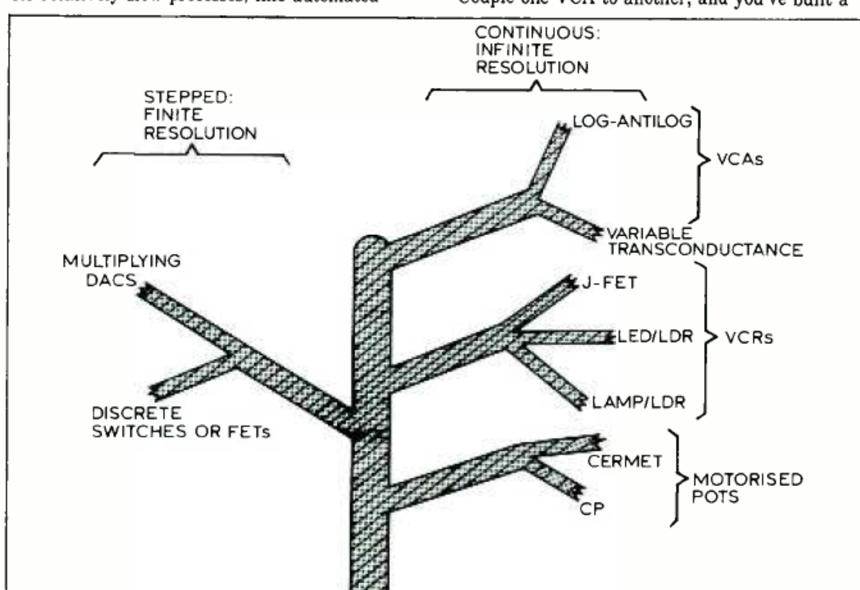

FIG 1: Tree of electronically-controlled gain and attenuation

root extractor or exponentiator. Add some non-linear elements (a diode or two) and you have a circuit which will 'solve' the complex and simultaneous non-linear equations of nature in real time.

In audiospeak, multiplication or division of a signal are synonymous with gain or attenuation, or decibels being added or subtracted. It's a blessing for simplicity that practical VCAs can embrace both directions or 'Quadrants' to a lesser or greater extent. It's also handy that the acronym VCA can refer to both Amplification and Attenuation; though at least one VCA manufacturer prefers to differentiate the latter mode, calling it *'VCAtt'*.

In its broadest sense, a VCA is a black box with three ports (**Fig 2**). Two are everyday (audio) input and output. A voltage (or current) applied to the third port, labelled 'CV' acts to alter the audio through-gain and hence the output level. To save on brain-damage, it's convenient if a linear change in control voltage (or current) causes an equally linear change in the dB level at the output, eg 10 dB per volt. Again, for many audio

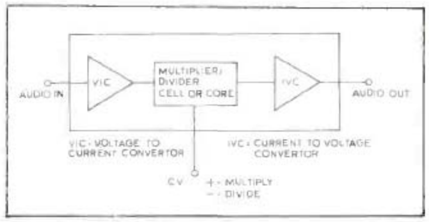

FIG 2: VCA generic block schematic

applications, the linearity of this relationship needs to be maintained over a range of some 3000x, or 70 dB. For others, 30 dB (30x) will suffice. Then for the kind of accuracy and repeatability that's desirable in up-market equipment, the ratio between DC control voltage and audio gain/attenuation will need to be tightly defined against the three Ts: temperature, time and manufacturing tolerance.

In the beginning, VCAs were built with tubes. The first solid-state circuits appeared in the early '60s. Using diodes and later FETs, they were sometimes called 'vario-lossers'. Flushed with the success of the first IC op-amps in the late '60s, IC makers specialising in analogue, notably Burr-Brown, the aptly named Analog Devices and National Semiconductor, began to develop at first hybrid, and then monolithic 'analogue multiplier/divider' chips, using bipolar transistors. Within a few years—by the mid '70s—some of their products were approaching the kind of performance that would be acceptable to pro-audio, but prices were high. Around the same time, a handful of specialist audio VCA makers arrived. Each appeared waving a patent which annexed one of the bridges that lay between a multiplier suited to monitoring processes in a sausage factory, and one that was good for fading audio. Today, the world of Audio VCAs revolves around two kinds of circuit topology, produced in volume by five US manufacturers. Both kinds are traceable to techniques that were first figured out for analogue computation about half a century ago.

Transconductance

The most rudimentary VCA one can build with transistors is just a simple development of the familiar differential amp, or 'long-tailed pair'. Looking into **Fig 3**, gain control is achieved by arranging for a voltage across 'CV' to vary the amount of (constant) current pulled through the amplifier transistors, TR1 and 2. The change in operating current directly affects their transconductance (current divided by voltage, or mA per volt=gm), hence the voltage gain, G recovered at A1's output (G=gm.RL). This kind of VCA and others whose operation relies on changing the ratio of voltage-to-current transfer of active devices, are loosely known as OTAs (Operational Transconductance Amplifiers), or as transconductance VCAs (or multipliers). National Semiconductor's LM 3080 is a well known example of the genre.

As it stands, the circuit has a limited range. For DC, a linear change in transconductance is advertised over 3 decades (60 dB), but for audio where the need for dynamic range enters the picture, non-linearity sets in early. For less than 1% distortion, the audio input is limited to a few millivolts, essentially because variations in transconductance can only produce changes in voltage gain if the circuit is operated open-loop, ie without NFB. Mechanisms include a progressive increase in the V_{BE} in TR.1 and 2 for tail currents above 1mA. The effect is only partially self-cancelling, hence the non-linearity. Distortion can be reduced for operation at normal line levels by passively attenuating the incoming audio (by around 40 dB), then recovering the level with an output amplifier. SNR suffers commensurately. Then, with high attenuation, slew rate becomes increasingly embarrassing as the active devices are starved of current. Instrumentation engineers refer to this circuit as a 'two quadrant multiplier'. This is to say that the audio or 'X input' has a bipolar capability (ie the signal can swing symmetrically), while the CV or 'Y input' is restricted to control voltages that are always positive (relative to the negative rail).

A more workable technique, employed in Nat Semi's LM 13700 (**Fig 4**), involves introducing a (nearly) constant bias current into each input node with a pair of diodes (D1,D2). Provided the diodes' geometry and temperature are similar to the devices in the OTA cell (Q4,5), their respective non-linearities are complementary, and partially cancel (**Fig 5**). With this *predistortion* method, drive levels can be raised by up to 15 dB ▷

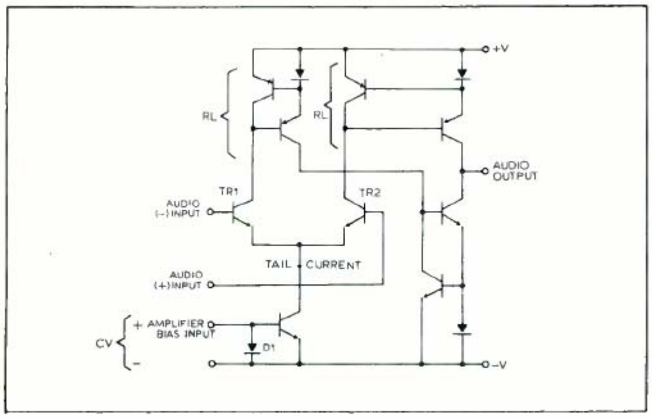

FIG 3: LM3080 operational transconductance amplifier

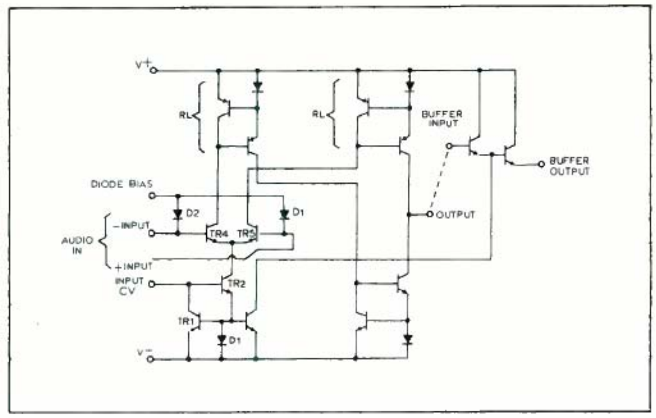

FIG 4: LM13700 operational transconductance amplifier with linearising diodes

for a given distortion level. SNR is improved in turn, but still remains around the level of cassette replay systems (circa −60 dB). And without input attenuation, fairly unnacceptable levels of THD persists for signals over 70 mV (−20 dBu). All told, NSC's 3080, 13700 and related chips from RCA were a godsend for the synthesisers and consumer-grade audio of the '70s and early '80s, but not much else.

In higher performance circuits, based on Barry Gilbert's classic 'current-ratioing' (or 'current-steering') transconductance cell[3], non-linearity is reduced by driving the active devices with current only. **Figs 6** and **7** illustrate typical circuits. The immediate trade-offs are added complexity, and the need for multiple, matched transistors (T1 to T4). The key is to keep the ratio of the currents in the cell's transistor-pairs constant and equal to the corresponding pair of external currents under all conditions. In effect, the cell transistors need to be dynamically matched, a multi-dimensional headache. The linear input range is expanded with a mixed bag of techniques like offsetting, pre-distortion networks, cross-coupling, base stoppers and emitter degeneration. Achieving low THD then hinges on the cell transistors' H_{fe} being kept as constant as possible over a scale of collector currents that's as broad as the desired control range, ie 10^5 for 100 dB. Subject to design finesse, transconductance cells of this genre can exhibit respectable audio specifications. They are particularly noted for wide bandwidth irrespective of gain or attenuation, and good isolation between audio (x) and the control signal (y).

In audio VCA parlance, good isolation is described as low 'control feedthrough'. The effect is a deviation from 0 v in the DC level at the VCA's output, usually referred to in millivolts. When defined with suitable reference to ZOL (ie

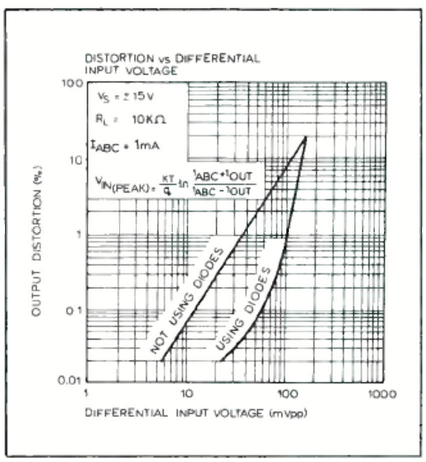

FIG 5: Operational transconductance amplifier with and without predistortion

as minus so many dBs) it helps the designer identify VCAs which don't go thump, pop, or click when the gain shift is rapid, which can be important in processors, and is vital for automated muting, but less so for automated faders. In the debit column, current-ratioing transconductance VCAs are primarily attenuators or dividers. They have to be tricked into producing gain, and in turn, their performance suffers. Ultimately, their clean operating range is constrained by imperfections in the cell's transistors, notably finite base-emitter resistance, V_{BE} mismatch and differences in the saturation current of individual transistors, as well as limitations in the topology, leading to inadequate common-mode gain at the input[4].

Log-antilog equation

Beyond variable transconductance cells, there are many other ways of implementing VCAs. Most of them are now obsolete, since audio isn't the only activity that can benefit from topologies that provide wide dynamic range, bandwidths beyond 20 kHz, and low distortion or 'error'. When these factors take precedence, there's only one other species in the race. The circuit in **Fig 8** exploits the almost perfect logarithmic relationship between a bi-polar transistor's base-emitter

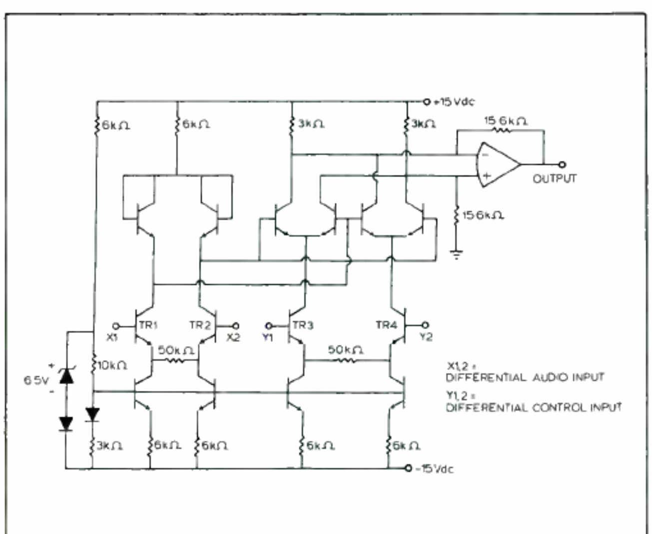

FIG 6: Four-quadrant variable transconductance multiplier

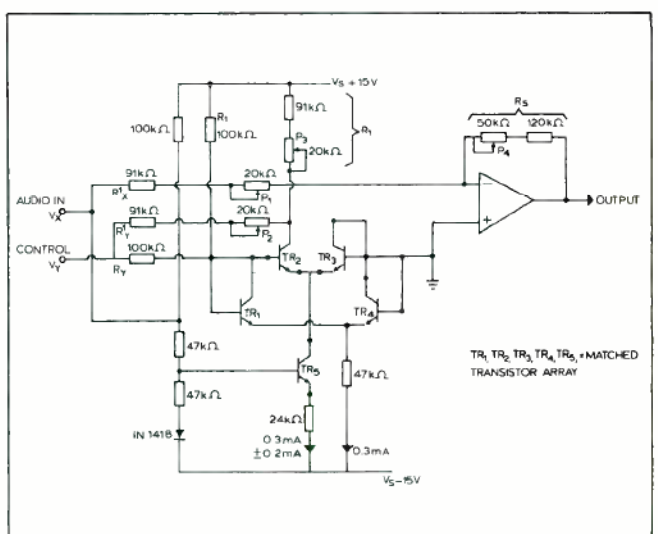

FIG 7: Four-quadrant variable transconductance multiplier

◁ voltage (V_{BE}) and collector current (I_C), plotted in **Fig 9**. Buffered by an op-amp, the combination acts as a 'log (arithmic) convertor'. A 60 mV change in V_{BE} (in TR.1) will result in, or be caused by (it all depends on topology): a tenfold change in I_C. Now the going is ratiometric, the same tenfold change can be described as a 20 dB change in the recovered output voltage. So the intrinsic scaling of the basic log circuit is 60 mV/20 dB or 3 mV per dB. The next piece of the jigsaw is illustrated in **Fig. 10**. Here, a logarithmic convertor has been combined with an anti-log convertor. With the control inputs V1 and V3 grounded, audio appearing at input V2 appears unchanged at the output. Sounds NBG, eh?

It all begins to make sense once it's recalled that adding and subtracting logs is equivalent to multiplication and division. And that with analogue electronics, addition and subtraction are child's play. When a control signal is applied to inputs V1 or V3, it's added to the logged or anti-logged signal emerging from T2 or T4, to produce a decrease or increase in gain respectively. Unity gain occurs when the audio input and output currents are equal, ie log of 1=0. The concise formula for the log-antilog VCA can now be written: (i) convert the audio and control voltages into logarithmic and linear currents respectively; (ii) combine them; (iii) anti-log the nett quantity; (iv) reconvert the signal into a voltage. The basic log-antilog core (T2, 4) is unipolar, ie operates in one quadrant only. The complementary control ports provide basic two quadrant operation. In other words, the basic log-antilog VCA offers both gain and attenuation, but as it stands, it can't handle bidirectional input signals, ie audio. As we'll see next month, this is one of the problems that log-antilog VCA pioneers Blackmer and Buff set out to conquer.

Compared to the current-ratioing transconductance species, the log-antilog VCA is adept at providing gain as well as attenuation. Another bonus is the control port's natural deci-linear relationship of 1 dB per 3 mV change in V_{BE}. In comparison, the control port law of transconductance VCAs is intrinsically exponential. However, as temperature changes, the log-antilog VCA develops a cumulative but entirely predictable error of +0.3% per +°C. Added to this, there's a highly variable error of 2.2 mV per °C, dependent on the temperature coefficient of the core transistors' base-emitter junctions, known as 'bulk offset voltage'. However, provided bulk offset is identical in all the core transistors, it will be cancelled by the reciprocal log-antilog action. Further, any deviations in the logarithmic relationship between V_{be} and I_c will result in harmonic distortion. Overall, low distortion depends on mutual matching of 2, 4 or 8 core transistors for log-conformance. Needless to say, the logarithmic characteristic gets progressively warped at extremes of the current scale. When the VCA gain is unity, mismatches in the log and anti-log transistors are self-cancelling. But with ascending gain or attenuation, distortion cancellation is progressively disorganised, and 'logging error' distortion reappears. The log-antilog VCA's control-feedthrough is potentially higher and the bandwidth, while ample for audio, does vary with gain setting.

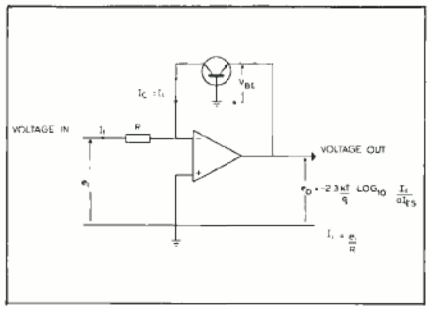

FIG 8: Log amplifier, transdiode configuration

In part 2, Ben Duncan describes how the circuit topologies just described are employed and refined in nine examples of modern VCA technology, including the products of the OEM VCA makers, whose chips populate the majority of pro-audio processors and consoles. □

Technical definitions and abbreviations

Cell Active heart of a variable transconductance VCA.
Core Active heart of a log-antilog VCA.
DNR Dynamic Range.
H_{FE} Current gain in a bipolar transistor.
I_C Collector current (in a bipolar transistor).
LDR Light Dependent Resistor (usually a Cadmium Sulphide cell).
NFB Negative Feedback.
SNR Signal to Noise Ratio.
SOTA State Of The Art.
V_{BE} Base-Emitter voltage of a bipolar transistor.
ZOL Zero Operating Level, eg. +4dBu.
H_{fe}, I_c, V_{be}: The lower case letters refer to incremental ('small signal') changes in these quantities.

References

1 J. Gordon Holt, Berning TF-12 preamplifier. *Stereophile*, July '88.
2 B. Duncan, Dynamic Loudness Compensation, Reproduced Sound 3, *Institute of Acoustics*, Nov '87.
3 B. Gilbert, A precise 4 Quadrant multiplier with sub-nanosecond response, *JSCC*, Vo.SC-3, Dec '68.
4 M. Hawksford, Low distortion programmable gain cell, using current-steering cascode topology, *J.AES*, Nov '82.

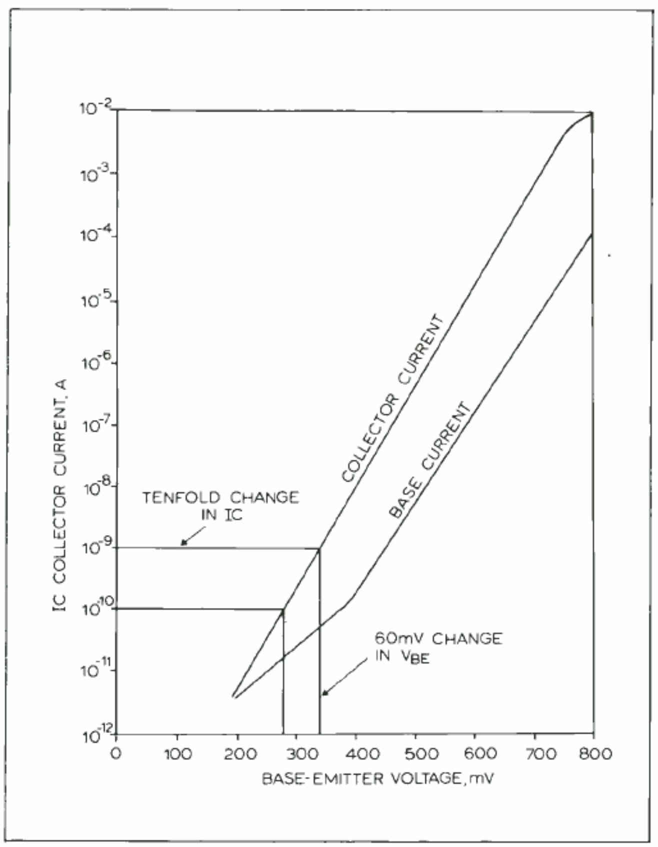

FIG 9: Collector and base current as a function of base-emitter forward bias with zero collector-base voltage

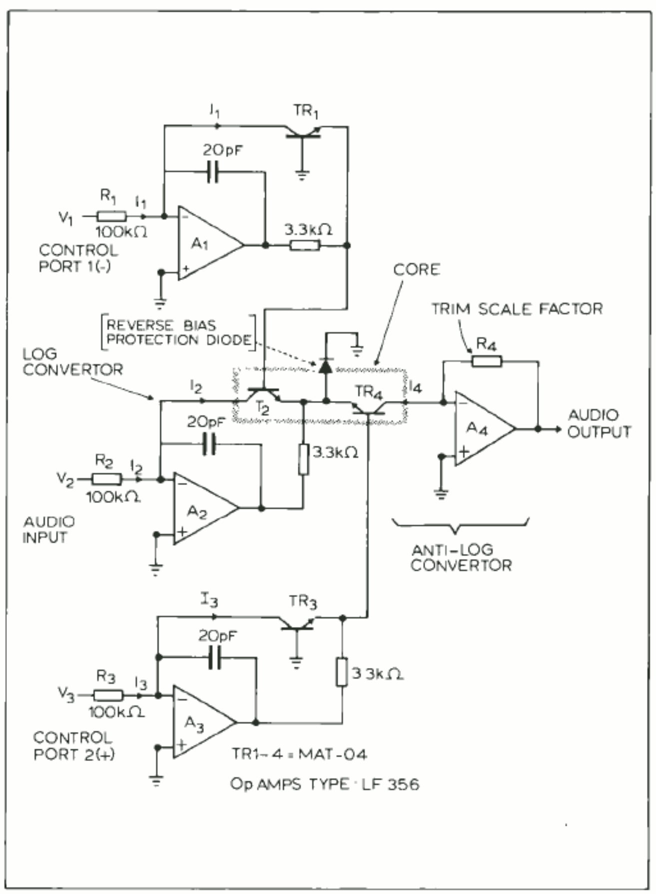

FIG 10: Log antilog multiplier/divider