Versatile analogue chip for oscilloscope plug-ins – Part I

In Part I of this two part feature *John Addis* of Tektronix Inc, Beaverton, Oregon, traces the history leading to the development of the M377, an 800MHz single chip amplifier used in the front end of the 11000 series of oscilloscopes.

The heart of four plug-in amplifiers for the Tektronix 11000 series oscilloscopes is a single high-speed bipolar analogue integrated circuit. The IC, internally designated the M377, forms almost the entire signal path in three of the plug-ins, and a majority of the fourth. With over 700 transistors, it is almost a "plug-in on a chip". A microprocessor and custom logic IC, control the M377 and add a sophisticated calibration routine.

The M377 implements:
● Gain switching over a 50:1 range in a 1, 2, 5 sequence of six discrete steps.
● A continuously variable gain control proportional to a dc voltage input.
● Multichannel operation with other M377s.
● Three identical outputs can each be independently inverted or enabled in less than 200ns.
● Bandwidth of 800MHz (420ps risetime) for gains of 0.4 to 12.
● 320MHz bandwidth at gain of 60.
● A four pole 100MHz bandwidth limit filter.
● A four pole 20MHz bandwidth limit filter.
● Differential input with high common mode rejection.
● Overdrive recovery to within 0.04% in 6ns.

To dissipate its three watts, the 4.32mm (0.170ins) × 2.92mm (0.115ins) M377 is mounted on a 1.22cm (0.480 inch) square thin film ceramic substrate. No bypass capacitors, resistors, or inductors are required, just a conductor pattern and the monolithic IC on a leadless substrate. The entire amplifier is connected to an etched circuit board using a new variation of Tektronix' patented Hypcon elastomeric connector.

Philosophy change

The M377's circuitry represents a radical architectural departure from earlier wideband oscilloscope amplifiers. Changing electronic industry needs for increased precision at high bandwidth and lower assembly costs dictated a fresh approach.

After the 7104 1GHz oscilloscope was introduced in 1979, some amplifier designers at Tektronix changed the focus of their interests. One gigahertz was just about the limiting system bandwidth obtainable with existing IC processes, and was enough to satisfy all but a few high frequency needs. There was, however, increased interest in high precision. The advent of 10 to 16 bit ADCs (analogue to digital converters), DACs, and digital signal processing required this precision. Other factors were also at work. Increased labour costs have made the assembly of complex, labour-intensive instruments too costly. Increasing the amount of circuitry in each IC should work as well in decreasing the assembly costs for analogue circuits as it has for digital circuits. Board space was another factor
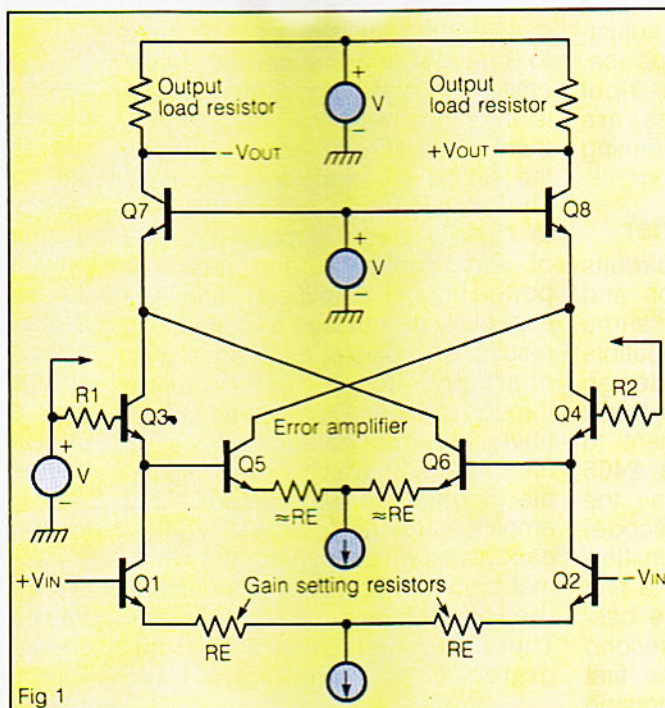

Figure 1: The compensated cascode "cascomp" amplifier

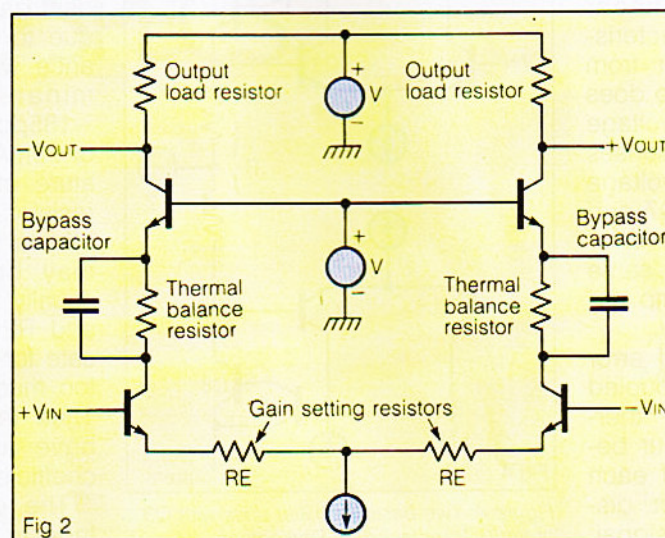Figure 2: Amplifier with thermal balance resistors and bypass capacitors

which dictated an increased level of integration. A four channel plug-in was required and there was simply not room for four channels of existing circuitry on the circuit board of one plug-in. Labour costs made calibration and adjustment expensive. Circuit designs that reduced the number of manual adjustments, thereby saving both labour costs and precious board space, were to be preferred.

The 11000 series oscilloscopes simplify dc calibration. The entire oscilloscope calibrates itself upon the user's command (or automatically). Only one manual adjustment per channel exists in most plug-ins using the M377. That one adjustment is for transient response. Even the high impedance input attenuator compensations are factory-adjusted by laser trimming under machine control.

The cascomp amplifier

It is possible to. design circuits which have greater precision and fewer adjustments at the same time. Indeed, these are compatible goals if the right circuit design choices are made. The Tektronix 2465, was our first instrument to attempt this approach. The 2465 relies upon a circuit we call the cascomp (compensated cascode) shown in figure 1. If the emitter-coupled differential pair was a *first generation* amplifier, then the cascomp can be considered a *second generation* amplifier. It is the first attempt at broadband precision amplification, and an improvement on the simple differential pair.

The cascomp senses the error voltage due to the non-linear (logarithmic) junction characteristics and subtracts the error from the final output. The cascomp does not sense the actual error voltage in Q1 and Q2, but instead senses the nearly identical error voltage generated by Q3 and Q4. Q7 and Q8 are required for good error amplifier bandwidth. For the same standing current, the cascomp demonstrates increased linearity.

There is another source of error in the first generation dc-coupled linear amplifier: "thermals" or "thermal tails". These effects occur because the operating point of each transistor and hence power dissipation change with the signal.
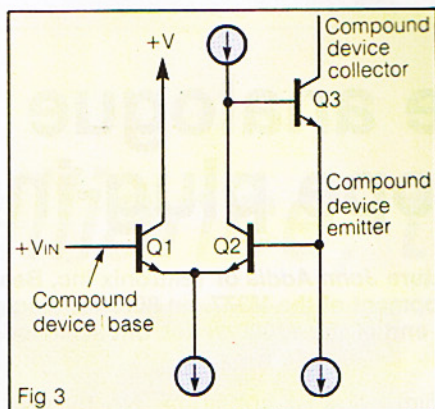

Figure 3: The basic broadband feedback amplifier

The power dissipation change causes the transistor's temperature to change and hence the base-emitter voltage. Since this voltage is in series with the input signal, there is a thermally induced error in the amplifier's transient response. The effect can be substantial (5% to 7% per stage in extreme cases of very wideband and therefore power hungry amplifiers), and it is especially troublesome because it results in multiple time constants, characteristic of heat diffusing through silicon. Discrete amplifiers have "thermal balance" resistors (see figure 2) which nearly eliminate thermals, but integrated circuit amplifiers cannot obtain the large capacitors which bypass the "thermal balance" resistors. Fortunately, the effects are usually linear. Therefore, first generation integrated circuit amplifiers have
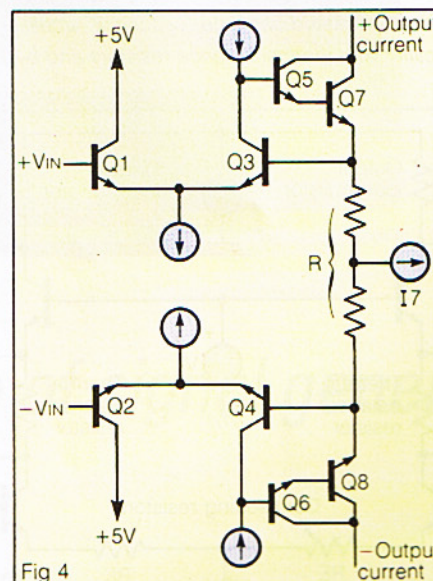

Figure 4: Two basic amplifier units with Q3 of figure 3 replaced by Darlingtons.

manual adjustments to compensate for signal-related thermal effects. For example, a 7104 oscilloscope and two single channel plug-ins require 32 manual adjustments to correct for thermals.

If the operating points of Q3 and Q4 (in figure 1) are chosen carefully, the cascomp cancels thermal errors as well as non-linear errors. By eliminating the need for manual adjustments, the parts cost, board area, and labour cost associated with calibration are reduced. The calibration labour costs are non-trivial because the adjustments for thermals are among the most difficult and frustrating to make.

The cascomp, a circuit patented by Pat Quinn of Tektronix, is successfully used in many Tektronix oscilloscopes today. It is an improvement over a simple differential pair, but has limitations. First, very high gain is not possible from a single stage without loss of all its advantages. Since the error amplifier and the main amplifier have the same gain, non-linearities in the error amplifier become appreciable at high gain. At very high gain, the error amplifier is not able to cancel non-linearities in a predictable manner.

Secondly, the stack of 3 transistor pairs in the signal path causes the gain to be 1.5 times as sensitive to alpha losses as the standard cascode. As temperature increases, the gain increases as α^3. This amounts to about 225ppm/°C in the cascomp for a typical β of 80 verses 150ppm/°C for a standard cascode. The cascode also has a counteracting temperature dependent gain term in the emitter circuit due to the dynamic emitter resistance which the cascomp has eliminated. This term is about −185ppm/°C for an emitter current of 20mA/side at a junction temperature of 60°C and a 40Ω emitter resistor. To compensate for alpha effects, two resistors, R1 and R2, may be added. However, it is usually not possible to make R1 and R2· large enough to compensate for alpha loss without creating too much high frequency peaking. Thus the standard cascode may have a lower gain temperature coefficient than a cascomp.

The cascomp has a disadvantage to IC design since the three

stacked devices use more of the available supply voltage. In circuits which require level shifting back down towards the negative supply, the voltage shift required is greater with cascomp than cascode.

A fourth disadvantage is in the cascomp's ability to handle overdrive signals. Since the error generating devices (Q3 and Q4) do not see the full input signal, the error correction circuitry and the main amplifier generally have different thermal histories in overdrive. With extra circuitry, thermals resulting from overdriving the cascomp can be about as good as a thermally balanced differential pair.

The basic amplifier

The M377 represents a third generation broadband amplifier design. The basic amplifier is shown in figure 3. This circuit, whose original design goes back at least to the mid seventies and is a variation on the LM102 of the late sixties, is a feedback amplifier. It can be viewed (imperfectly) as a compound transistor in which a differential pair, Q1 and Q2, compares an input signal with the emitter voltage of an output device, Q3. The compound device has increased gm and β over Q3 alone operating at the same current.

When viewed as a compound transistor, it is obvious that the output can be taken from either the emitter or the collector of Q3. The analogy with a compound transistor falls apart because the α of Q3 is not helped by Q1 and Q2. This flaw can be corrected in several ways. For example, the collector of Q1 could be connected to the emitter of Q3. This not only increases the α of the compound device, it also bootstraps Q1's collector. However, the operating point of Q1 compromises its Ft. Additionally, at very high frequencies this design has a negative input impedance and is potentially unstable. The method most suited to high frequency designs is to change Q3 into a Darling-

Figure 5: Method of varying broadband voltage amplification using a resistor network and transistor selection

ton as in figure 4 and accept the lost advantages of bootstrapping as a price to pay for stability and a 1GHz bandwidth.

When two such voltage followers are connected together as in figure 4, the configuration is that of an instrumentation amplifier, the basic amplifier stage used in the M377. A block diagram of the M377 would show only two stages of gain. In between are a level shift, a Gilbert multiplier variable gain control, and a choice of two bandwidth limit filters of a full bandwidth path. The input stage and the three identical

output stages are all instrumentation amplifiers whose inputs are the bases of Q1 and Q2 and whose outputs are the Darlington collectors. Gain for these stages is set by the resistance R between mirror-imaged voltage followers. To high precision, the signal current output is equal to the voltage input between Q1 and Q2 bases divided by R.

Electronic gain switching

The ability to change gain over a wide range in precision steps is important in oscilloscope preamplifiers. It allows the input attenuator to be simplified to just three settings, unity, ten times attenuation, and one hundred times attenuation. This is important because the input attenuator, which is built with electro-mechanical relays, is inherently expensive and less reliable than well-designed solid state components. The fewer the components, the more reliable the attenuator will be. Changing gain also results in the best trade-off between gain and noise. For example, if the amplifier needs a maximum sensitivity of 1mV per division, use of a passive attenuator before the amplifier as the *only* means of altering the amplification would result in the high gain amplifier's full noise level being displayed at all sensitivities. Use of a passive attenuator after amplification increases the dynamic range requirements for any amplifier before the attenuator. Even at 10:1 attenuation, the cost of increased dynamic range can be severe.

There are several ways of changing the gain in an IC. One method is to use a passive constant resistance network as shown in figure 5. Current in the collector of Q1 will result in different values of Vout depending upon whether Q2, Q3, or Q4 is conducting. This method can have very broad bandwidth and has the advantage of minimal change in bandwidth between attenuations selected. Used alone, it has the disadvantage of
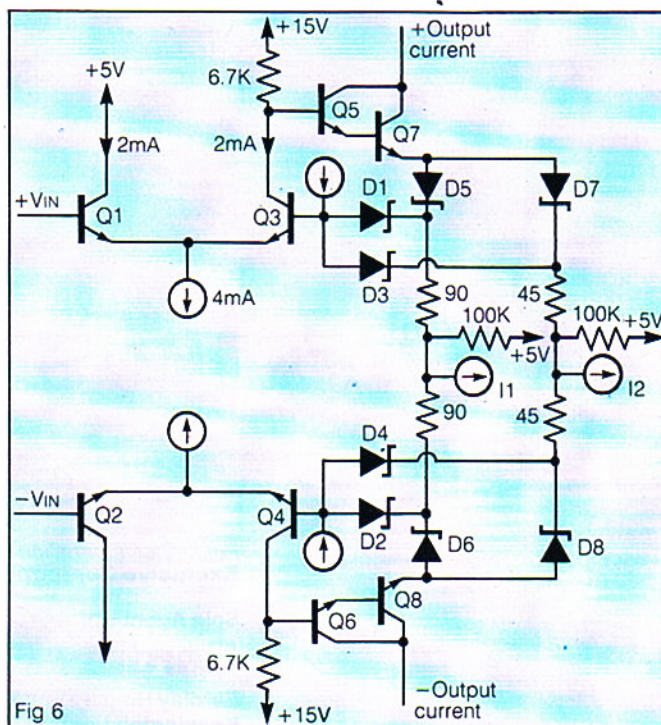
Figure 6: Two stages of variable gain amplifier, sources I, and I$_2$ force standing current into input stages.

Instrumentation amplifier

requiring the amplifier preceding it to handle the full dynamic range of the lowest gain setting. The input amplifier cannot simultaneously have high gain and wide dynamic range without prohibitive power dissipation, so it must have low gain to handle the largest expected signals. This requires the post attenuator amplifiers to operate at high gain with consequently high output noise. Some means of controlling the input amplifier's gain offers a way around this dilemma. For example, increasing the input transistor's emitter resistor will allow the amplifier to handle large signals at low gain settings. Decreasing the emitter resistor provides enough gain to minimise noise from subsequent stages at high gain settings. The disadvantage of this approach is that the bandwidth will change from one gain setting to another. In the M377, discrete gain steps are obtained by changing the first stage gain setting resistor (as shown in figure 6). TTL logic inputs select one of six current sources such as I1 or I2 which force the first stage standing current through the appropriate gain setting resistors and associated diodes. (Only two gain settings are shown in figure 6.) Diodes D1 and D2 or D3 and D4 close the feedback loop, and deselected diodes are back biased by the appropriate 100kΩ resistor. The diodes must withstand the brunt of the full input signal swing. Since the diode-connected high frequency transistor will punch through at about 1.5V reverse bias, Schottky barrier diodes which can withstand reasonable voltages are used. Two 6.7kΩ nichrome resistors are used in place of active current sources to reduce noise.

In the linear range, thermal effects are quite small in this amplifier when compared to a simple differential pair because the large standing current and current swings necessary to produce an output signal are handled by Q7 and Q8. Thermal effects in these devices are inside the feedback loop and are reduced about 250 times by the voltage gain of Q1 and Q3 or Q2 and Q4. The operating points of Q1, Q2, Q3, and Q4 are virtually unaffected by input signals. The voltage change on these devices is small because the input signal is small and the current change is small because Q5, Q6, Q7, and Q8 have such high current gain. In fact, Q1, Q2, Q3, and Q4 have barely changed their operating points at all as the signal cuts off either Q7 or Q8. Tight thermal coupling of Q1, Q2, Q3, and Q4 keep the remaining linear (as opposed to overdrive) thermals well below the 0.05% level. However, once Q7 or Q8 cuts off, the four input devices rapidly overdrive trying to bring the amplifier back into the linear range. Taming the resulting overdrive thermals is a unique accomplishment of the M377.

Part II of the development of this third generation instrumentation front end amplifier continues in the September 1988 edition with more detail of the design of the M377 and the effect it had on other aspects of oscilloscope design and the resulting solutions.□

Versatile analogue chip for oscilloscope plug-ins Part II

In Part II of this two part feature, *John Addis* of Tektronix Inc, continues with a discussion of more detail, and some of the implications that the performance of the M377, an 800MHz single chip amplifier has on a number of aspects of oscilloscope design.

I was aware that feedback amplifiers such as the 7A22 plug-in had less linear and overdrive thermal effects than some of the Tektronix high speed amplifiers. It was the small thermal efects and improved linarity which made me first choose a feedback amplifier as the M377's basic design, but it was not until late in the design that a way of radically improving the overdrive recovery occurred.

The M377 is capable of recovering from overdrive signals of up to 2 volts to within 0.04% in about 6ns. Recovery to the 0.01% level takes 25ns. This is more than three orders of magnitude faster recovery to the same level than the 7A13, a plugin designed some years ago specifically for fast recovery. It is about 2 orders of magnitude faster in setting than good modern operational amplifiers, and about five times faster than the best 12 bit DACs.

An awareness of just how good 0.02% recovery in 20ns is became apparent as we tried to make the rest of the plug-in signal path support that kind of performance from the amplifier. For example, a 42 inch length of RG 58 coaxial cable, terminated in 50ohms at each end still exhibits 0.02% skin effect loss 200ns after a transient even though the total transit time of the cable is only 5ns! The small diameter cable just 10 inches long used to connect the front panel BNC connector to the M377 inside the 11A52 plug-in exhibits 0.02%

skin effect loss of its own about 20ns after a step, for the first time these losses became easy to measure and provided us with more than a few surprises. For example, the small diameter cable paradoxically settled to its final value much more rapidly than larger RG-58 of
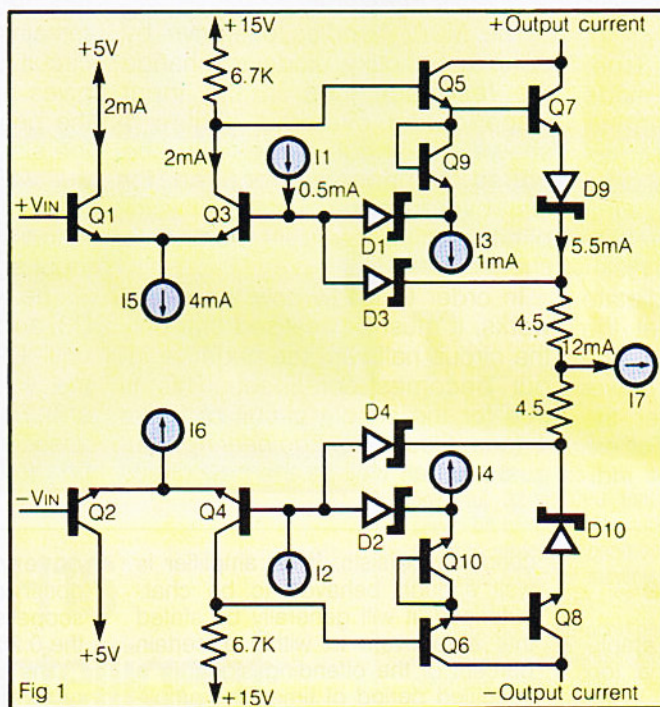

Fig 1

Figure 1: Amplifier design with overdrive recovery components

the same delay. As we thought about our measurements, the reason became apparent. The small cable has a copperweld centre conductor of copper plated steel. The dc loss of this cable is much greater per foot than the RG 58, and the time required to settle to its final (dc) loss is therefore less. An alternative explanation in the frequency domain: The dc loss equals the skin effect loss at a much higher crossover frequency.

Another discovery was a little more painful. The 11A52 plug-in has only a 50ohm input impedance. This provides greater bandwidth and lower noise than a plug-in which requires a buffer amplifier with a 1 megohm input impedance. In fact, two M377s, one for each channel, are all the amplification the 11A52 has. The 11A52 then requires a 50ohm attenuator which must be switched via relays (the 11A52, as with the entire 11000 series is completely programmable). We chose a high reliability version of the popular subminiature style hermetically sealed relay. These relays require glass to metal seals for hermeticity and use kovar for leads and header because kovar matches the thermal expansion coefficient of glass. But kovar is twenty to fifty times as resistive as copper and is ferromagnetic as well. Kovar is a great candidate for skin effect loss. Most of the kovar is gold plated, but the small segment inside the glass seal is not because the gold must be plated after the very high temperature glass to metal seal is made. It is in this very short section where serious skin effect losses occur. We found that, although the total path length through 5 relays in the attenuator was only 4 inches, the skin effect error did not die out to the 0.02% level for a full microsecond! To further complicate matters, when the relays are driven from a high impedance, such as a passive probe, the skin effect

almost disappears. The reason for this is that the skin effect loss is an increase in the series path resistance at high frequencies. When the attenuator is driven by a high impedance, the series loss is swamped out by the source's resistance.

Our solution was to introduce a network of resistors and capacitors in the attenuator which compensated for the skin effect loss. The whole attenuator has a 9% dc loss, but is flat from dc to 1GHz within about 2%. The compensated attenuator is flat to within 0.05% at 40ns, compared to about 2% at 40ns before compensation. The attenuator risetime is just 130ps.

Overdrive recovery in the IC

Precision overdrive recovery is a subject of some myth and misunderstanding. Most designers think about preventing transistors from saturating when they think of quick overdrive recovery. True, Schottky TTL switches are made faster than TTL by preventing saturation. But amplifiers are not the same as digital circuits, and transistor saturation is not what prevents a transistor amplifier from recovering quickly to within 0.1%.

It is thermal effects which dictate the amplifier recovery time at the (typically) 0.5% level and less. The faster the amplifier, the most power it dissipates, and the greater the potential for thermal effects. For example, amplifiers in the 7104 indi-

Fig 2
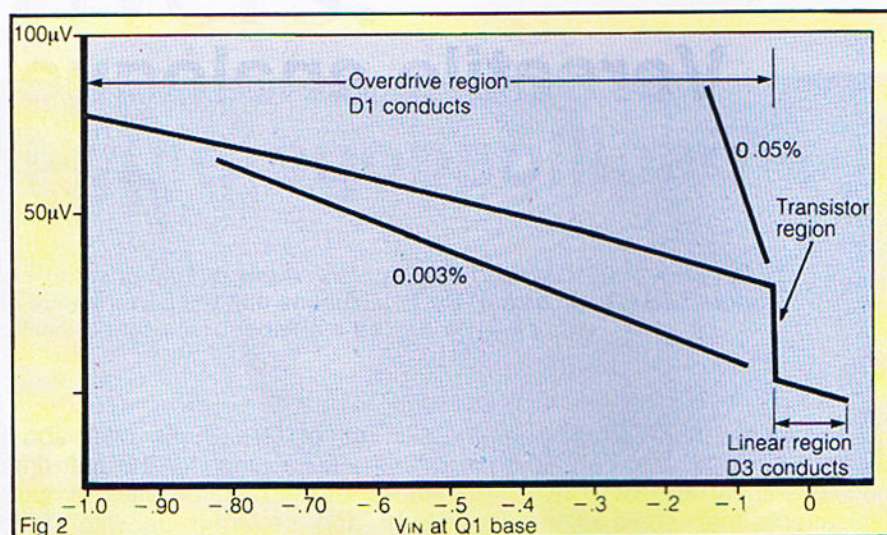
Figure 2: Thermally induced error voltages

vidually have a bandwidth of about 2.5GHz per stage and thermals of about 6% per stage.

The M377 handles overdrive by using a Schottky diode to change the feedback loop in the input stage during overdrive. Figure 1 shows the input stage with the added components required for fast overdrive recovery. For simplicity, only one gain setting is shown.

In order to follow how the circuit works, it must be realised that only the circuit half with the negative input becomes non-linear. This is true for the simple circuit of figure 4 (see *Electronic Engineering*, August pxx) as well as the new circuit

of figure 1. For positive input signals, the amplifier takes all the current in the current source I7 and remains linear, but the mirror image circuit becomes non-linear when it gives up all its standing current to the positive input side. It is always the side with the more negative input which goes non-linear. Therefore it is only necessary to understand what happens for negative input signals.

The voltage at the cathodes of D3 and D9 will follow +Vin down until D9 and Q7 cut off because the voltage at the junction of D4 and D10 is held by −Vin and its associated amplifier. At that point, D1 conducts because Q5 and Q9

are still conducting due to current source I3. D1 steals the current flowing through D3 and cuts D3 off. D1 closes a new feedback loop consisting of Q3, Q5, Q9, and D1. Without any loading, this changed circuit will follow the input signal down until I5 saturates or Q5 breaks down.

Figure 2 shows the thermals generated in Q1 and Q3 over the input voltage range of −1V to 0V. Q2 and Q4 generate no appreciable thermals when a positive signal is applied to Q1. The linear input range is about ±50mV and the gain is very high due to the 9ohm total emitter load. Because of tight thermal coupling, the thermal voltages generated are only 80μV per milliwatt of power difference between Q1 and Q3.

Variable gain control

Almost every oscilloscope has a continuously variable gain control. This control is not as frequently used as the coarse (step) control, yet it is a source of several design difficulties. Passive attenuation (potentiometer) works well up to about 50MHz. The chief advantage of this scheme is low cost and the chief disadvantages are the limited bandwidth and mechanical constraints. Above 50MHz some form of electronic gain control based upon the Gilbert multiplier is usually used. While the Gilbert multiplier can work well, it is not without its imperfections. Chief among these are the low bandwidth when compared to other stages in the same IC, the addition of thermal effects, and the generation of noise.

Automatic calibration placed requirements upon the M377 design, requirements which ultimately became useful features. One of these requirements is precision electronic gain control. In order for a microprocessor to calibrate plug-in gain, the gain control element must be predictable and stable. At the same time, this precision allowed calibrated deflection factors in between the coarse step attenuation settings. Thus the plug-in deflection factor is always calibrated to better than 1% accuracy, even at 4.51mV/div for example.

There are several Gilbert multiplier configurations to choose from. Figure 3 shows the usual four-


Figure 3: A four quadrant Gilbert multiplier

quadrant multiplier. This configuration allows the input signal in the form of a current in I1 and its complement I2 to be multiplied from +1 though zero to −1 as it appears in the output collectors. But there is a problem with this configuration. As desired, the low frequency signal current splits between emitters in the same ratio as the dc standing current does, but at high frequencies the signal splits in accordance with the emitter impedances at that frequency. That impedance is determined primarily by r_b at the highest frequencies where beta is low. Since the four devices are equal in size, the gain at the highest frequencies tends to be zero. In fact, there are only three gains for which the frequency response is theoretically flat, +1, 0, and −1. By symmetry, it is apparent that +1, 0, and −1 are also the only gains for which thermals are theoretically zero. The M377 multiplier is used at gains between +1 and +0.3, so flat frequency response at a gain of zero is of no benefit.

The four-quadrant multiplier's noise as a function of signal gain is lowest at −1 and +1, but maximum at zero gain. This is because any base resistance generates a thermal noise which is amplified by both emitter coupled pairs and added in the output. When the mul-

tiplier's signal gain is zero, both emitter coupled pairs have maximum gain and therefore maximum noise. The S/N ratio is zero! At signal gains of +1 and −1, the voltage gain for noise generated in base resistors is zero.

Another Gilbert multiplier configuration is of greater use to the M377, and is shown in figure 4. Here the current in the inner pair of resistors, Q2 and Q3, is wasted while the signal is taken strictly from the outer pair, Q1 and Q4. Since currents in the inner pair match the currents in the outer pair when the signal gain is 0.5, thermals are zero and high frequency current split is perfect by symmetry. Furthermore, half the noise current is thrown away with the unwanted signal current, and in addition the noise is less! Although the frequency response and thermals are not perfect at gains other than 0, 0.5, and 1, the gain never strays so far from these ideal points as it does in the four-quadrant multiplier. Furthermore, the frequency response is improved because the two-quadrant multiplier in figure 4 has only one collector's worth of capacitance on each output instead of two.

A possible disadvantage is that the two-quadrant multiplier will not allow inversion, while the four-quadrant multiplier will. In practice

though, there are ways of inverting a signal much more accurately than the Gilbert multiplier. For greater versatility, the M377 employs a separate inverter for each of its three outputs. The signal inversion is accomplished with less than 0.02% gain change.

Band gap reference

The variable gain control is controlled by an analogue input voltage between −1.0V (zero gain) and +1.0V (full gain). Except at 0V input where the gain is one half that set by the coarse gain control, the gain must be referenced to an absolute dc voltage. For this reason, an internal band gap reference was added. With the band gap, changes in power supply due to load changes in the instrument could not affect the gain accuracy.

Another advantage of an internal reference is ruggedness. An internal band gap can have zero output until the +5V supply is at least +3V. All the M377 current sources from the −5V supply are referenced to the band gap. Those current sources cannot turn on in the absence of a +5V supply, and therefore no transistor can go into saturation or cause latch-up if the +5V supply is lost. Of course, the current sources from the −5V supply cannot be on without the −5V supply either, hence there is no power supply sequencing necessary and loss of any supply safely shuts down the chip.

Multi-channel operation

The M377 was originally designed for the 11A32, 11A33, 11A34, and 11A52 plug-in amplifiers.

The 11A32 and 11A52 are two channel plug-ins, the 11A33 is a single channel plug-in, and the 11A34 is a 4 channel plug-in. The M377 is a single channel amplifier with three separate outputs, one for display, one for trigger, and an additional channel for auxiliary purposes. Each output's impedance and common mode voltage level remain the same whether an output is enabled or not, so outputs from two or more different M377s can be connected in parallel to form a channel switch. Any output may be selected, the outputs can be alternated or chopped, and the add or sum mode is accomplished by
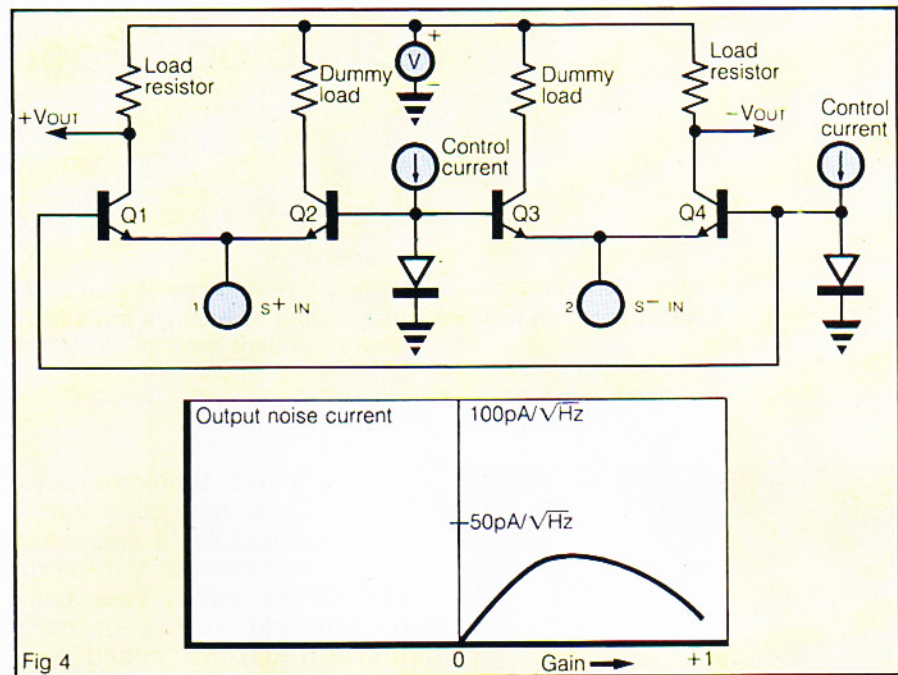

Figure 4: A two quadrant Gilbert multiplier

turning on two channels. Because each output can be inverted in only 200ns, it is even possible to display the sum and difference of two channels with chopped or alternate sweep mode. Since each M377 output is separately controlled, the trigger mode can be different than the display mode.

To accomplish the same gain at the output of each plug-in independent of the number of channels, the M377 is laser trimmed while still in the wafer stage to have an output impedance of 50Ω, 100Ω, or 200Ω per side. The two channel plug-ins use the 100Ω version and the four channel plug-in uses the 200Ω part. In this way, the overall plug-in output impedance is 50Ω per side and the gain is the same independent of the number of channels.

In the case of two channel plug-ins, 100Ω transmission lines are used to connect together the outputs from two M377s. The combined output is taken midway between the two M377s. This results in a nominally perfect 50Ω output impedance at all frequencies.

The four channel plug-in cannot be similarly reverse terminated because it is not possible to construct 200Ω transmission lines on etched circuit boards. Reflections among the four chips cancel, so the small loss of bandwidth which occurs is due primarily to the increased capacitive loading on the output due to the chips themselves. Reverse termination is compromised by the 100Ω, 125ps, transmission lines used to connect the 200Ω outputs together. Since the oscilloscope mainframe is terminated in 50Ω, there is little signal reflected back to the plug-ins.

The ability to laser trim nichrome resistors while the M377 is still in wafer form affords the opportunity to trim common mode output level to zero. Gain at the six discrete steps is trimmed to 1% tolerance and two gain settings are trimmed for dc balance as well. The variable gain control (Gilbert multiplier) is trimmed so that full gain and zero gain occur with +1.0 volts and −1.0 volts at the analog gain control input respectively. In all, eighteen resistors are trimmed and several qualifying tests are done in about 60s.

Finally, on the trimmed wafer, 126 dc tests are performed, some of them to a 0.02% test limit, in 11s per chip.

Acknowledgements