8. One Trip Down the IC Development Road

This is the story of the last IC that I developed. I use the word develop rather than design because there is so much more involved in the making of a standard part than just the circuit design and layout. My goal is to give the reader an idea of what is involved in this total development. The majority of this description will be on the evolution of the product definition and the circuit design since that is my major responsibility. I will also describe many of the other important steps that are part of the IC development. To give the reader an idea of what is required, I made an approximate list of the steps involved in the development of an IC.

The steps in the development of a new IC:

1. Definition
2. Circuit design
3. Re-definition
4. More circuit design
5. The first finalizing of the specifications
6. Test system definition
7. Mask design
8. Test system design
9. Waiting for wafers to be made
10. Evaluation
11. Test system debug
12. Redesign (circuit & masks)
13. More waiting
14. Finalizing the test system
15. IC characterization
16. Setting the real specifications
17. Pricing
18. Writing the data sheet
19. Promotion
20. Yield enhancements

Circuit design (steps 2, 4, and 12) is what we usually think of when we talk about IC design. As you can see, it is only a small part of the IC development. At some companies, particularly those that do custom ICs, circuit design is all the design engineers do. In the ideal world of some MBAs, the customer does the definition, the designer makes the IC, the

test engineer tests, the market sets the price, and life is a breeze. This simple approach rarely develops an IC that is really new; and the companies that work this way rarely make any money selling ICs.

Most successful IC designers I know are very good circuit designers and enjoy circuit design more than anything else at work. But it is not just their circuit design skills that make these designers successful; it is also their realization that all the steps in the development of an IC must be done properly. These designers do not work to a rigid set of specifications. They learn and understand what the IC specs mean to the customer and how the IC specs affect the system performance. Successful IC designers take the time to do whatever it takes to make the best IC they can.

This is quite different from the custom IC designer who sells design. If you are selling design, it is a disadvantage to beat the customer's spec by too much. If you do the job too well, the customer will not need a new custom IC very soon. But if you just meet the requirement, then in only a year or so the customer will be back for more. This kind of design reminds me of the famous Russian weight lifter who set many world records. For many years he was able to break his own world record by lifting only a fraction of a kilogram more than the last time. He received a bonus every time he set a new world record; his job was setting records. He would be out of a job if he did the best he could every time; so he only did as much as was required.

Product Definition

Where do we get the ideas for new products? From our customers, of course. It is not easy, however. Most customers will tell you what they want, because they are not sure what they need. Also, they do not know what the different IC technologies are capable of and what trade-offs must be made to improve various areas of performance. The way questions are asked often determines the answers. Never say, "Would you like feature XYZ?" Instead say, "What would feature XYZ be worth to you?"

When an IC manufacturer asks a customer, it is often like a grandparent asking a grandchild. The child wants all the things that it cannot get from its parents and knows none of the restrictions that bind the others. The only thing worse would be to have a total stranger do the questioning. That may sound unlikely, but there are companies that have hired non-technical people to ask customers what new products they want. At best, this only results in a very humorous presentation that wastes a lot of people's time.

Talking to customers, applications engineers, and salespeople gives the clues and ideas to a designer for what products will be successful. It is important to pick a product based on the market it will serve. Do not make a new IC because the circuit design is fun or easy. Remember that circuit design is only a small part of the development process. The days of designing a new function that has no specific market should be long

gone. Although I have seen some products recently that appear to be solutions looking for problems!

This is not to say that you need marketing surveys with lots of paperwork and calculations on a spreadsheet. These things are often management methods to define responsibility and place blame. It is my experience that the errors in these forms are always in the estimate of the selling price and the size of the market. These inputs usually come from marketing and maybe that is why there is such a high turnover of personnel in semiconductor marketing departments. After all, if the marketers who made the estimates change jobs every three years, no one will ever catch up with them. This is because it typically takes two years for development and two more years to see if the product meets its sales goals.

So with almost no official marketing input, but based on conversations with many people over several years, I began the definition of a new product. I felt there was a market for an IC video fader and that the market was going to grow significantly over the next five years. The driving force behind this growth would be PC based multi-media systems. At the same time I recognized that a fader with only one input driven is a very good adjustable gain amplifier and that is a very versatile analog building block. The main source of this market information was conversations with customers trying to use a transconductance amplifier that I had designed several years earlier in fader and gain control applications.

The Video Fader

The first step is figuring out what a video fader is. The basic fader circuit has two signal inputs, a control input and one output. A block diagram of a fader is shown in Figure 8–1. The control signal varies the gain of the two inputs such that at one extreme the output is all one input and at the other extreme it is the other input. The control is linear; i.e., for the control signal at 50%, the output is the sum of one half of input 1 and one half of input 2. If both inputs are the same, the output is independent of the control signal. Of course implementing the controlled potentiometer is the challenging part of the circuit design.

The circuit must have flat response (0.1dB) from DC to 5MHz and low differential gain and phase (0.1% & 0.1 degree) for composite video applications. For computer RGB applications the -3dB bandwidth must be at least 30MHz and the gain accuracy between parts should be better than 3%. The IC should operate on supply voltages from ±5V to ±15V, since there are still a lot of systems today on ±12V even though the trend is to ±5V. Of course if the circuit could operate on a single +5V supply, that would be ideal for the PC based multi-media market.

The control input can be in many forms. Zero to one or ten volts is common as are bipolar signals around zero. Some systems use current inputs or resistors into the summing node of an op amp. In variable gain amplifier applications often several control inputs are summed together.

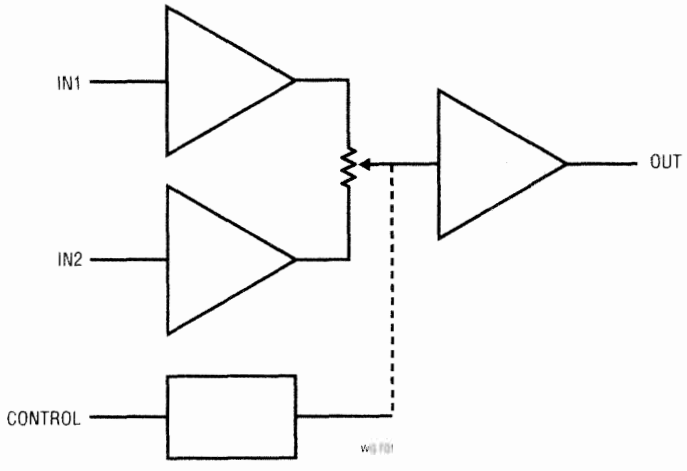

Figure 8–1.
Basic fader circuit.

In order to make a standard IC that is compatible with as many systems as possible, it is desirable to make the control input user defined. At the same time it is important that the IC not require a lot of external parts.

To make the circuit more immune to errors in the potentiometer circuit, we can take feedback from the output back to both inputs. Figure 8–2 shows this feedback and replaces the potentiometer with the mathematical equivalent blocks: K, 1-K, and summation. Now the output is better controlled, since the value of K does not determine the total gain, only the ratio of the two input signals at the output. The gain is set by the feedback resistors and, to a smaller degree, the openloop gain of the amplifiers.

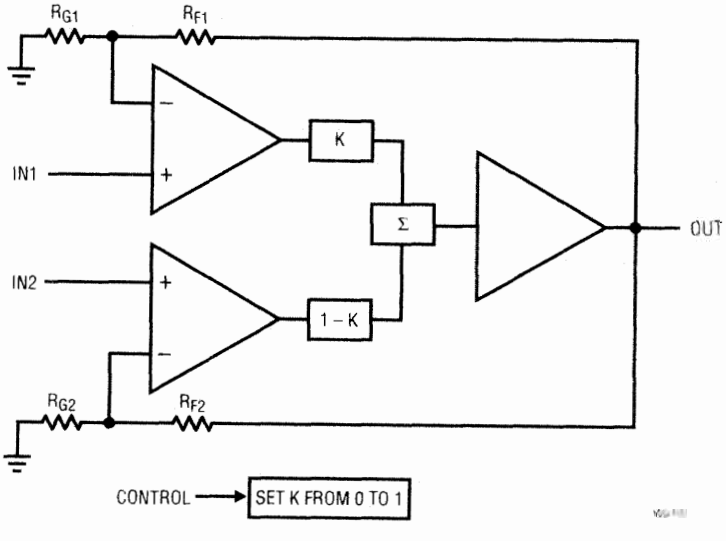

Figure 8–2.
Feedback fader circuit.

Circuits

At this point it is time to look at some actual circuits. Do we use voltage feedback or current feedback? Since the current feedback topology has inherently better linearity and transient response, it seemed a natural for the input stages. One customer showed me a class A, current feedback circuit being implemented with discrete transistors. Figure 8–3 shows the basic circuit. For the moment we will not concern ourselves with how the control signal, V_C, is generated to drive the current steering pairs. Notice that the fader is operating inverting; for AC signals this is not usually a problem, but video signals are uni-polar and another inversion would eventually be needed. I assumed that the inverting topology was chosen to reduce the amount of distortion generated by the bias resistors, R_{B1} and R_{B2}, in the input stages.

Since transistors are smaller than resistors in an IC, I intended to replace the bias resistors with current sources. Therefore my circuit could operate non-inverting as well as inverting, and as a bonus the circuit would have good supply rejection. The complementary bipolar process that I planned to use would make class AB implementations fairly straightforward. I began my circuit simulations with the circuit of Figure 8–4; notice that there are twice as many components compared to the discrete circuit and it is operating non-inverting.

After a bit of tweaking the feedback resistor values and the compensation capacitor, the circuit worked quite well. The transistor sizes and

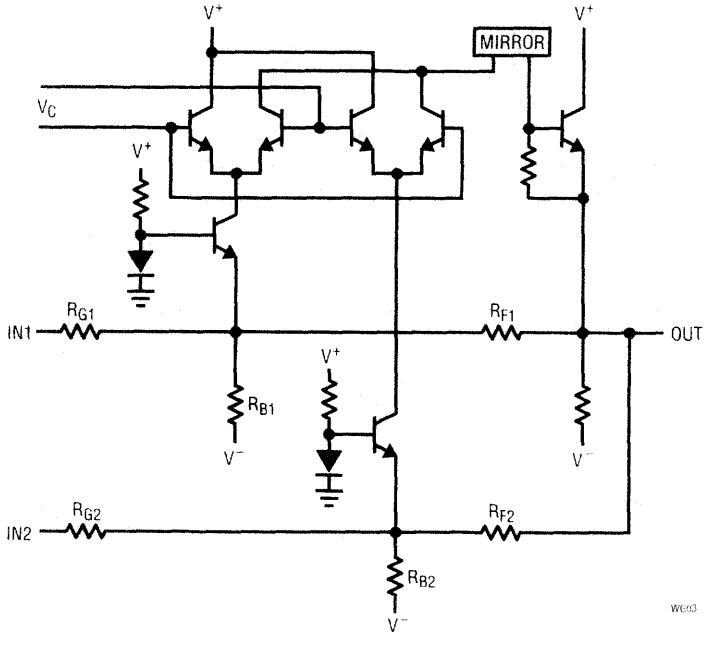

Figure 8–3. Discrete design, class A current feedback fader.

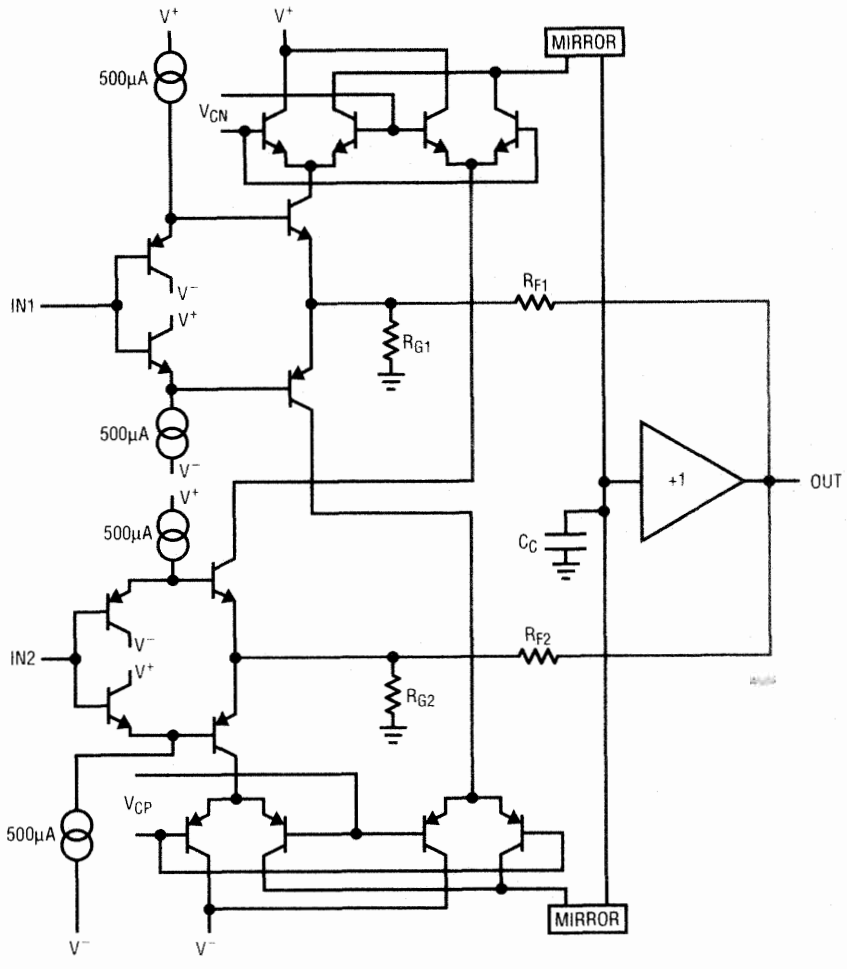

Figure 8–4.
Class AB current feedback fader.

current levels were set based on previous current feedback amplifiers already designed. It was time to proceed to the control section.

For linear control of the currents being steered by a differential pair, the voltage at the bases of the steering transistors must have a nonlinear characteristic. This TANH characteristic is easily generated with "pre-distortion" diodes. The only requirement is that the currents feeding the diodes must be in the same ratio as the currents to be steered. The circuit of Figure 8–5 takes two input control currents, K and (1-K), and uses Q1 and Q2 as the pre-distortion diodes to generate the control signal V_{CN} for the NPN steering transistors. The collector currents of Q1 and Q2 then feed the pre-distortion diodes Q3 and Q4 that generate V_{CP} to control the PNP steering transistors.

I noticed that the linearity of the signal gain versus diode current is strongly influenced by the bulk R_b and R_e of the current steering transistors. After consulting some papers on multipliers (thank you Barry Gilbert) I found that there are some topologies where the bulk R_b and R_e of the pre-distortion diodes compensate the equivalent in the steering

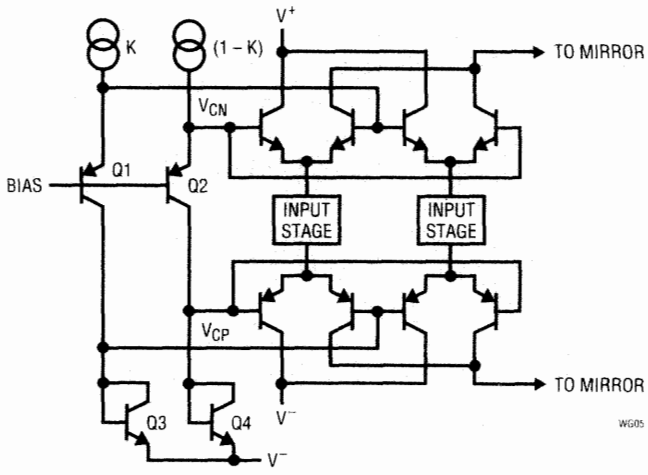

Figure 8–5. Basic circuit to drive the steering transistors.

transistors. Unfortunately, in my circuit I am using PNPs to drive NPNs and vice versa. In order to match the pre-distortion diodes to the steering transistors, a more complicated circuit was required. I spent a little time and added a lot more transistors to come up with a circuit where the pre-distortion diodes for the NPN steering transistors were NPNs, and the same for the PNPs. Imagine my surprise when it didn't solve the linearity problem. I have not included this circuit because I don't remember it; after all, it didn't work.

So I had to learn a little more about how my circuit really worked. In the fader circuit, the DC current ratio in the steering transistors is not important; the small signal current steering sets the ratio of the two inputs. Figure 8–6 shows a simplified circuit of the pre-distortion diodes and the steering transistors. The diodes and transistors are assumed perfect with 18Ω resistors in series with the emitters to represent the bulk R_b and R_e of the devices. The control currents are at a 10:1 ratio; the DC currents in the

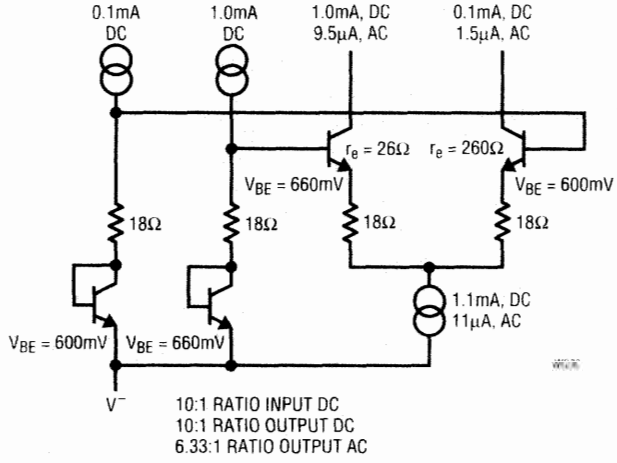

Figure 8–6. Bulk resistance problems in steering.

steering transistors are also at a 10:1 ratio. But the small signal steering is set by the ratio of the sum of the r_e and the bulk resistance in each transistor, and in this case the result is a 6.33:1 ratio!

In the fader circuit, the only way to improve the gain accuracy is with low R_b and R_e steering transistors. Unfortunately this requires larger transistors running at low current densities and that significantly reduces the speed (F-tau) of the current steering devices. I went back to the simpler circuit of Figure 8–5, increased the size of the current steering transistors, and tweaked the compensation capacitor and feedback resistors to optimize the response.

Now it was time to find a way to interface the external control signal(s) to the pre-distortion diodes of Figure 8–5. The incoming signal would have to be converted to a current to drive the pre-distortion diodes, Q1 and Q3. A replica of that current would have to be subtracted from a fixed DC current and the result would drive the other pre-distortion diodes, Q2 and Q4.

I did not want to include an absolute reference in this product for several reasons. An internal reference would have to be available for the external control circuitry to use, in order not to increase the errors caused by multiple references. Therefore it would have to be capable of significant output drive and tolerant of unusual loading. In short, the internal reference would have to be as good as a standard reference. The inaccuracy of an internal reference would add to the part-to-part variations unless it was trimmed to a very accurate value. Both of these requirements would increase the die size and/or the pin count of the IC. Lastly, there is no standard for the incoming signals, so what value should the reference be?

I decided to require that an external reference, or "full scale" voltage, would be applied to the part. With an external full scale and control voltage, I could use identical circuits to convert the two voltages into two currents. The value of the full scale voltage is not critical because only the ratio between it and the control voltage matters. With the same circuit being used for both converters, the ratio matching should be excellent.

Figure 8–7 shows the basic block diagram that I generated to determine what currents would be needed in the control section. The gain control accuracy requirements dictated that an open loop voltage-to-current converter would be unacceptable. Therefore a simple op amp with feedback would be necessary. It became clear that two control currents (I_C) were needed but only one full scale current (I_{FS}) was. Mirror #1 must have an accurate gain of unity in order to generate the proper difference signal for mirror #3. Mirrors 2 and 3 must match well, but their absolute accuracy is not important. All three mirrors must operate from zero to full scale current and therefore cannot have resistive degeneration that could change their gain with current level.

In order to use identical circuits for both voltage-to-current converters, I decided to generate two full scale currents and use the extra one to bias the rest of the amplifiers. You can never have too many bias currents available.

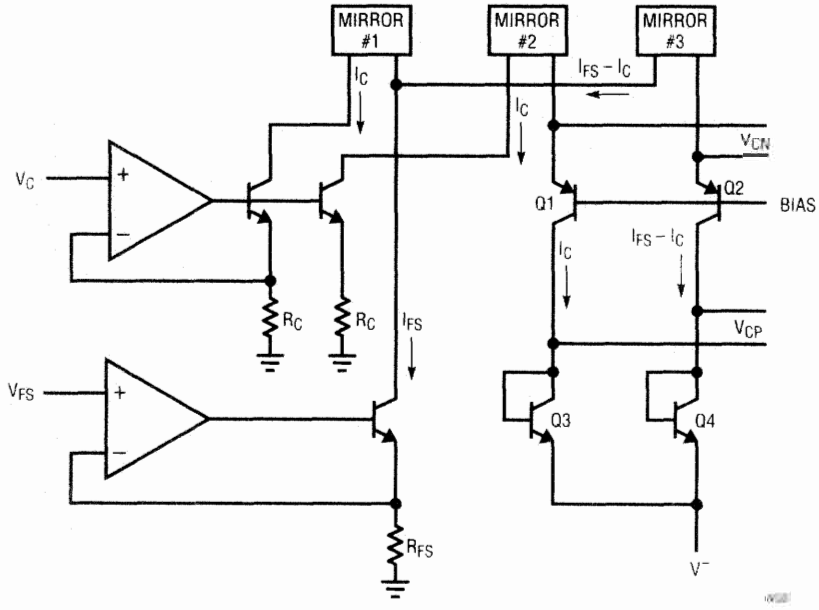

Figure 8–7.
Block diagram of the control circuit.

The block diagram of Figure 8–7 became the circuit of Figure 8–8 after several iterations. The common mode range of the simple op amp includes the negative supply and the circuit has sufficient gain for the job. Small current sharing resistors, R1, R2, R3, and R4, were added to improve the high current matching of the two output currents and eliminate the need for the two R_C resistors. The small resistors were scaled so they could be used for short circuit protection with Q5 and Q6 as well.

Mirror #l is a "super diode" connection that reduces base current errors by beta; the diode matches the collector emitter voltages of the matched transistors. Identical mirrors were used for #2 and #3 so that any errors would ratio out. Since these mirrors feed the emitters of the pre-distortion cascodes Q1 and Q2, their output impedance is not critical and they are not cascoded. This allows the bias voltage at the base of Q1 and Q2 to be only two diode drops below the supply, maximizing the common mode range of the input stages.

While evaluating the full circuit, I noticed that when one input was supposed to be off, its input signal would leak through to the output. The level increased with frequency, as though it was due to capacitive feedthrough. The beauty of SPICE came in handy now. I replaced the current steering transistors with ideal devices and still had the problem. Slowly I came to the realization that the feedthrough at the output was coming from the feedback resistor. In a current feedback amplifier, the inverting input is driven from the non-inverting input by a buffer amp and therefore the input signal is always present at the inverting input. Therefore the amount of signal at the output is just the ratio of the feedback resistor to the amplifier output impedance. Of course the output impedance rises with frequency because of the single pole compensation necessary to keep

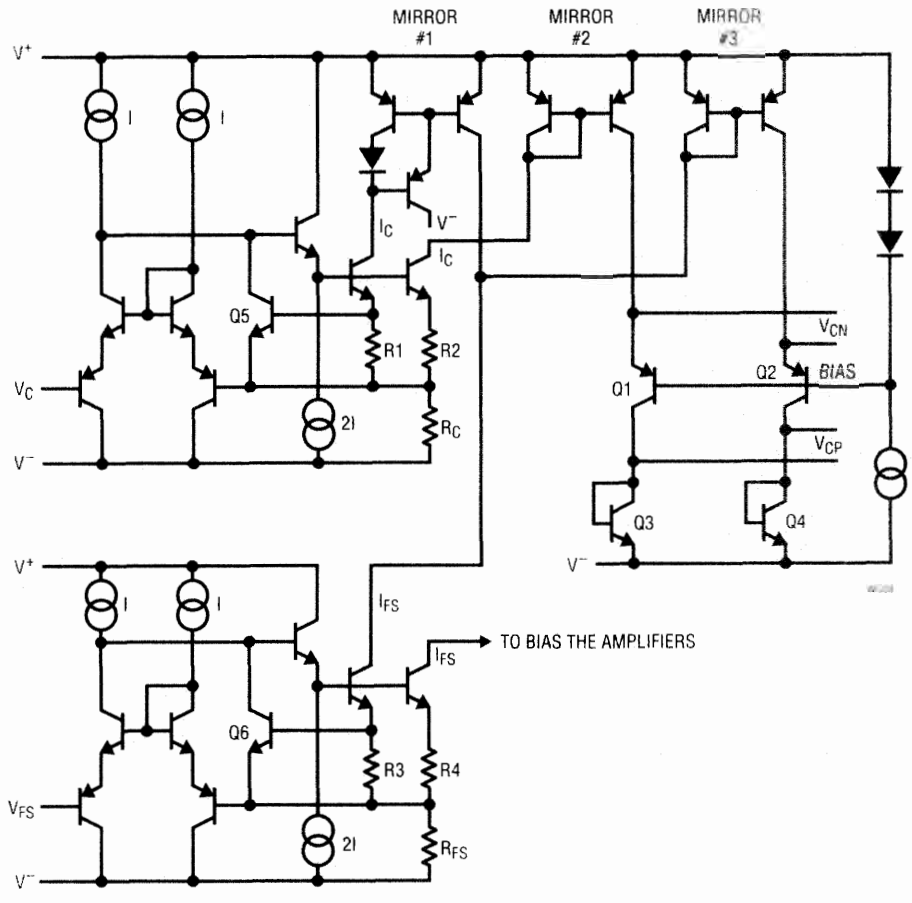

Figure 8–8.
The control circuit.

the amplifier stable. The basic current feedback topology I had chosen was the feedthrough problem. Now it was obvious why the discrete circuit was operating inverting. The problem goes away when the non-inverting input is grounded because then the inverting input has very little signal on it.

Redefinition

At this point I realized I must go back to the beginning and look at voltage feedback. I started with the basic folded cascode topology and sketched out the circuit of Figure 8–9. It seemed to work and there were no feedthrough problems. It also appeared to simplify the control requirements, since there were no PNPs to steer. While working with this circuit I realized that the folded cascode transistors, Q7 and Q8, could be used as the steering devices, and sketched out Figure 8–10. This looked great since it had fewer devices in the signal path and therefore better bandwidth. The only downside I could see was the critical matching of the current sources; all eight current sources are involved in setting the gain. While I was pondering how to get eight current sources coming

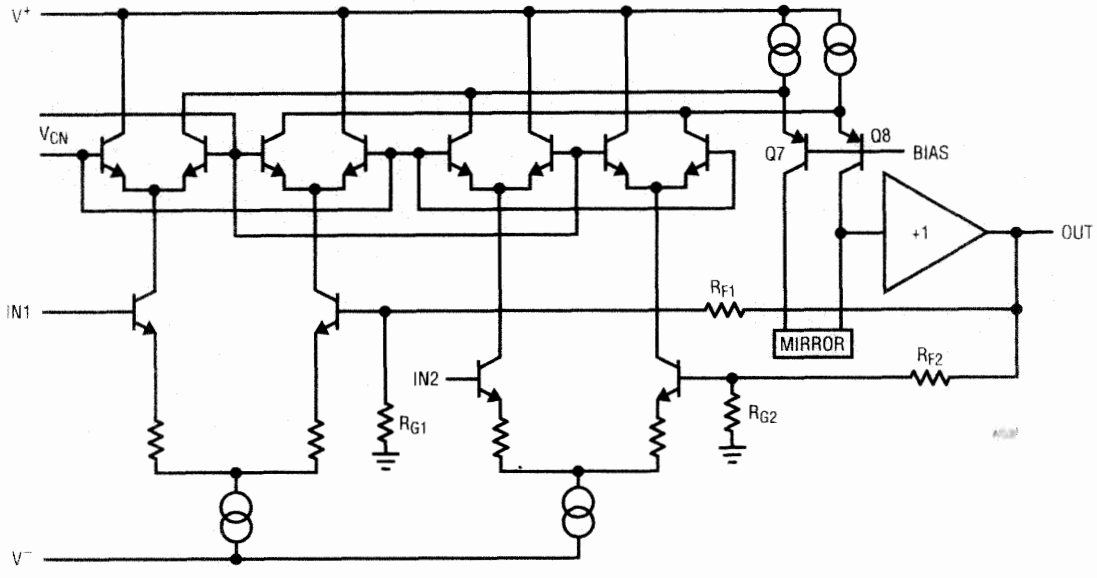

Figure 8–9.
Voltage feedback fader.

from opposite supplies to match, I decided to run a transient response to determine how much input degeneration was required.

The bottom fell out! When the fader is set for 10% output, the differential input voltage is 90% of the input signal! This means that the *open loop* linearity of the input stage must be very good for signals up to one volt or more. To get signal linearity of 0.1% would require over a volt of degeneration. With that much degeneration in each input stage, the mismatch in offset voltage between the two would be tens of millivolts and that would show up as control feedthrough. Big degeneration resistors

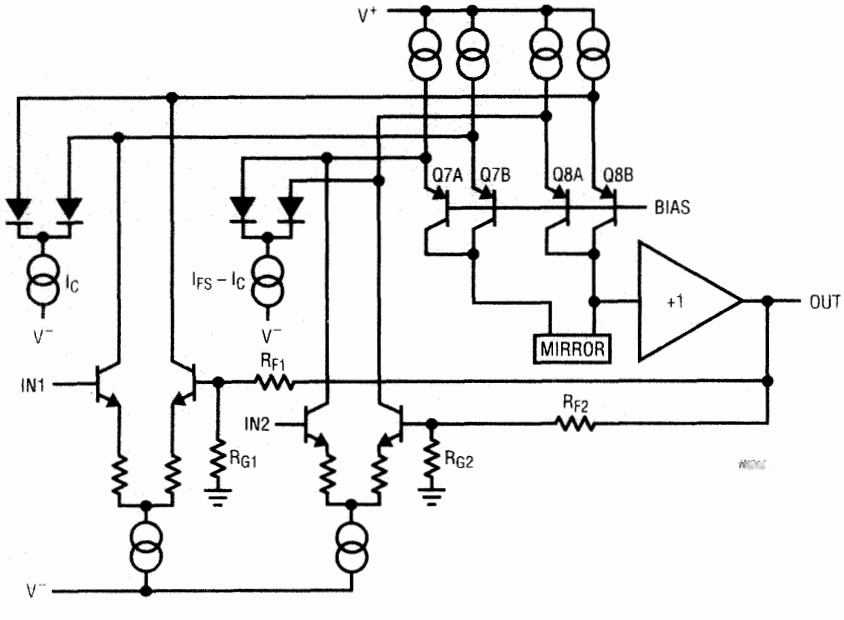

Figure 8–10.
Voltage feedback with cascode steering.

also generate serious noise problems and cause the tail pole to move in, reducing the speed of the amplifier. It was time to retreat to the current feedback approach and see how good I could make it.

The current feedback topology has very low feedthrough when operated inverting, so I started with that approach. Unfortunately the feedthrough was not as good as I expected and I started looking for the cause. The source of feedthrough was found to be the emitter-base capacitance of the current steering transistor coupling signal into the pre-distortion diode that was holding the transistor off. Unfortunately the off diode was high impedance (no current in it) so the signal then coupled through the collector base capacitance of the steering transistor into the collector, where it was not supposed to be. Since the steering transistors had to be large for low R_b and R_e, the only way to eliminate this problem was to lower the impedance at the bases of the steering transistors.

What I needed was four buffer amplifiers between each of the four pre-distortion diodes and the current steering transistors. To preserve the pre-distortion diodes' accuracy, the input bias current of the buffers needed to be less than one microamp. The offset of the buffers had to be less than a diode drop in order to preserve the input stage common mode range so that the circuit would work on a single 5V supply. Lastly, the output impedance should be as low as possible to minimize the feedthrough.

The first buffer I tried was a cascode of two emitter followers, as shown on the left in Figure 8–11. By varying the currents in the followers and looking at the overall circuit feedthrough, I determined that the output impedance of the buffers needed to be less than 75Ω for an acceptable feedthrough performance of 60dB at 5MHz. I then tried several closed loop buffers to see if I could lower the supply current. The circuit shown in Figure 8–11 did the job and saved about 200 microamps of supply current per buffer. The closed loop buffer has an output impedance of about 7Ω that rises to 65Ω at 5MHz. Since four buffers were required, the supply current reduction of 800 microamps was significant.

At this point it became obvious to me that for the feedthrough to be down 60dB or more, the control circuitry had to be very accurate. If the full scale voltage was 2.5V and the control voltage was 0V, the offset errors had to be less than 2.5mV for 60dB of off isolation. Even if I trimmed the IC to zero offset, the system accuracy requirement was still very tough. I therefore wanted to come up with a circuit that would insure that the correct input was on and the other input was fully off when the control was close to zero or full scale. I thought about adding intentional offset voltage and/or gain errors to the V-to-I converters to get this result, but it didn't feel good. What was needed was an internal circuit that would sense when the control was below 5% or above 95% and force the pre-distortion diodes to 0% and 100%. Since the diodes were fed with currents, it seemed that sensing current was the way to go.

Since the currents that feed the pre-distortion diodes come from identical mirrors, I wanted to see if I could modify the mirrors so that they

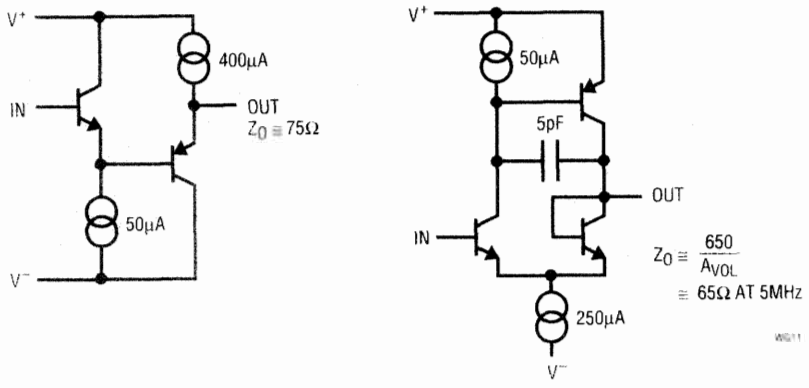

Figure 8–11. Open- and closed-loop buffers.

would turn off at low currents. This would work at both ends of the control signal because one mirror is always headed towards zero current. The first thought was to put in a small fixed current that subtracted from the input current. This would add an offset near zero (good) and a gain error everywhere else (bad). Now if I could turn off the offset current when the output current was on, it would be perfect. Current mirrors #2 and #3 in Figure 8–8 were each modified to be as shown in Figure 8–12. The offset current is generated by Q9. A small ratio of the output current is used to turn off Q9 by raising its emitter. The ratios are set such that the output goes to zero with the input at about 5% of full scale. The nice thing about this mirror is that the turn-off circuit has no effect on mirror accuracy for inputs of 10% or more. The diode was added to equalize the collector-base voltage of all the matching transistors.

At this point the circuit was working very well in the inverting mode and I went back to non-inverting to see how the feedthrough looked. Since the output impedance of the amplifier determines the feedthrough performance, I eliminated all the output stage degeneration resistors. I set the output quiescent current at 2.5 milliamps so the output devices would be well up on their F-tau curve and the open loop output impedance would be well under 10 Ohms. The feedthrough was still 60dB down at 5MHz. I added a current limit circuit that sensed the output transistors' collector current, and the circuit topology was finalized.

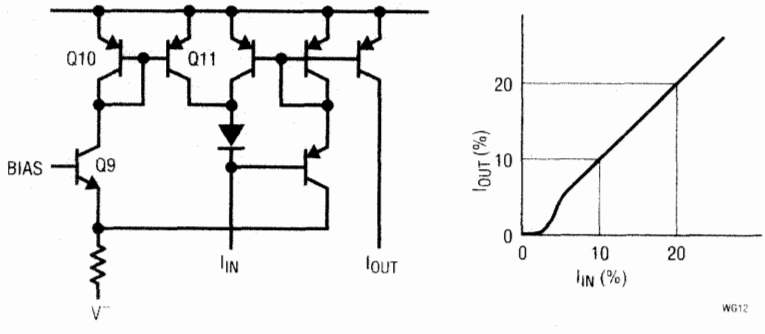

Figure 8–12. Mirror with a turn-off.

The last step in the circuit design is rechecking and/or optimizing the area of every transistor. This is usually done by checking the circuit's performance over temperature. I always add a little extra area to the transistors that are running close to saturation when the additional parasitic capacitance won't hurt anything.

Mask Design

Experienced analog IC designers know how important IC layout is. Transistors that are supposed to match must have the same emitter size and orientation as well as the same temperature. The fader output amplifier is capable of driving a cable and generating significant thermal gradients in the IC. For this reason I put both input stages on one end of the die next to the current steering devices and put the output stage at the other end. The bias circuits and the control op amps went in the middle. The best way to minimize thermal feedback is distance. The 14-pin SO package set the maximum die size and the pad locations.

The IC process used had only one layer of metalization and therefore I provided the mask designer with an estimate of where "cross-unders" would be needed. For those of you not familiar with the term "cross-under," I will explain. A cross-under is a small resistor, usually made of N+, inserted in a lead so that it can "cross-under" another metal trace. Normally these cross-unders are inserted in the collectors of transistors, since a little extra resistance in the collector has minimal effect.

The fader circuit, with over 140 transistors and very few resistors, was clearly going to have a lot of cross-unders. I was resigned that both supplies would have many cross-unders; in order for the circuit to work properly, the voltage drops introduced by the cross-unders must not disturb the circuit. For example, the current mirrors will common mode out any variation in supply voltage as long as all the emitters are at the same voltage. This is easy to do if the emitters all connect together on one trace and then that trace connects to the supply. As mask design progresses, it is important that each cross-under added to the layout be added to the schematic and that circuit simulation is re-checked. Time spent before the silicon comes out to insure that the circuit works is well spent.

I would like to make a comment or two on mask design and the time that it takes. For as long as I can remember, speeding up mask design has been the Holy Grail. Many, including myself, have thought that some new tool or technique will cut the time required to layout an IC significantly. When computer layout tools became available, they were sold as a productivity enhancement that would cut the time it takes to layout ICs. The reality was that the ICs became more complex and the time stayed about the same.

The analog ASIC concept of a huge library of functions available as standard cells that are just plopped down and hooked up sounds great; except that very few innovative products can be done with standard func-

tions. What typically happens is that each new product requires modifications to the "standard" cells or needs some new standard cells. You're right back at transistor level optimizing the IC. Of course no one ever plans for the extra time that this transistor level optimization takes, so the project gets behind schedule.

The "mono-chip" or "master-chip" idea is often used to speed up development. This technique uses just the metal layer(s) to make the new product; a large standard IC with many transistors and resistors is the common base. The trade-off for time saved in mask design is a larger die size. The argument is often made that if the product is successful, a full re-layout can be done to reduce die size and costs. Of course, this would then require all the effort that should have been done in the first place. I would not argue to save time and money up front because I did not expect my part to be successful!

In summary, mask design is a critical part of analog IC development and must be considered as important as any other step. Doing a poor job of mask design will hurt performance and that will impact the success of a product much more than the extra time in development.

Testing

IC automatic test system development is an art that combines analog hardware and software programming. We cannot sell performance that we cannot test. It is much easier to measure IC performance on the bench than in an automatic handler. In successful companies, the good test development engineers are well respected.

The fader IC requires that the closed loop AC gain be measured very accurately. The gain is trimmed at wafer sort by adjusting the value of resistor R_C. This trim is done with the control input fixed and the linearity of the circuit determines the gain accuracy elsewhere. The errors due to the bulk resistance of the steering transistors have no effect at 50% gain; therefore it seemed like the best place to trim the gain.

While characterizing the parts from the first wafer, I noticed that there were a few parts that had more error than I expected at 90% gain. I also determined that these parts would be fine if I had trimmed them at 90%. It was also true that the parts that were fine at 90% would not suffer from being trimmed at 90%. So, I changed my mind as to where the circuit was to be trimmed and the test engineer modified the sort program. More wafers were sorted and full characterization began.

Setting the data sheet limits is a laborious process that seems like it should be simpler. The designer and product engineer go over the distribution plots from each test to determine the maximum and minimum limits. In a perfect world we would have the full process spread represented in these distributions. Even with a "design of experiments" run that should give us the full spread of process variations, we will come up short of information. It's Murphy's law. This is where the designer's knowledge

of which specs are important, and which are not, comes into play. It makes no sense to "over spec" a parameter that the customer is not concerned about because later it could cause a yield problem. On the other hand, it is important to spec all parameters so that any "sports" (oddball parts) are eliminated, since they are usually caused by defects and will often act strangely. The idea is to have all functional parts meet spec if they are normal.

Data Sheets

The data sheet is the most important sales tool the sales people have. Therefore it is important that the data sheet is clear and accurate. A good data sheet is always late. I say this based on empirical data, but there seems to be a logical explanation. The data sheet is useless unless it has all the minimums and maximums that guarantee IC performance; as soon as those numbers are known, the part is ready to sell and we need the data sheet. Of course it takes time to generate the artwork and print the data sheet and so it is late. One solution to this problem is to put out an early, but incomplete, data sheet and then follow it a few months later with a final, complete one.

Analog ICs usually operate over a wide range of conditions and the typical curves in the data sheet are often used to estimate the IC performance under conditions different from those described in the electrical table. The generation of these curves is time consuming and, when done well, requires a fair amount of thought. Human nature being what it is, most people would rather read a table than a graph, even though a table is just an abbreviated version of the data. As a result, the same information is often found in several places within the data sheet. I am often amazed at how inconsistent some data sheets are; just for fun, compare the data on the front page with the electrical tables and the graphs.

Beware of typical specs that are much better than the minimums and maximums. I once worked with a design engineer who argued that the typical value should be the average of the distribution; he insisted that the typical offset voltage of his part was zero even though the limits were ±4mV. Most companies have informal definitions of "typical", and it often varies from department to department. George Erdi added a note to several dual op amp data sheets defining the typical value as the value that would yield 60% based on the distributions of the individual amplifiers. I like and use this definition but obviously not everyone does, since I often see typicals that are 20 times better than the limits! Occasionally the limits are based on automatic testing restrictions and the typicals are real; for example, CMOS logic input leakage current is less than a few nanoamps, but the resolution of the test system sets the limit at 1 microamp.

Summary

Since you are still reading, I hope this long-winded trip was worth it. The development of an IC is fun and challenging. I spent most of this article describing the circuit design because I like circuit design. I hope, however, that I have made it clear how important the other parts of the development process are. There are still more phases of development that I have not mentioned; pricing, press releases, advertising, and applications support are all part of a successful new product development. At the time of this writing, the video fader had not yet reached these phases. Since I am not always accurate at describing the future, I will not even try. Those of you who want to know more about the fader should see the LT1251 data sheet.

At this time I would like to thank all of the people who made the video fader a reality and especially Julie Brown for mask design, Jim Sousae for characterization, Dung (Zoom) Nguyen for test development, and Judd Murkland in product engineering. It takes a team to make things happen and this is an excellent one.