PROVOCATIVE TOPICS ON LINEAR INTEGRATED CIRCUITS

Robert A. Pease
Staff Scientist,
National Semiconductor Corp.
Mail Stop C2500A
2900 Semiconductor Drive
Santa Clara CA 95051

Abstract

A number of provocative topics will be discussed such as: relative priorities, Check Lists, Design Reviews, Design Rule Checks, digitally computerized simulation, analog computer simulations, breadboards, test strategies, and metal-mask options. These are demonstrated to be valuable for enhancing the probabilities of good results in high-volume linear Integrated Circuits (ICs).

Historical Considerations

It is well known that the design of linear integrated circuits has been, historically, an arcane art, accomplished by bizarre persons with strange personalities. While there may be an element of truth to this, it is also fair to say that reasonable, rational, and sane persons can also design successful ICs. But for best results, they should take up some of the more fanatical attitudes that the old-timers used to have. A mere enthusiasm is not necessarily sufficient for good results, but a meticulous, skeptical, suspicious approach sometimes is helpful to find and avoid possible problems that could wreck or delay the project. This is especially true for "large-volume" linear ICs where it has been observed that a few "minor" problems can turn the project into a "low-volume" one. I shall try to present examples and approaches which have been found useful to avoid or solve the problems.

Note that I am dwelling primarily on linear ICs. This is partly because the design of large digital ICs has often turned into a committee effort, with many man-years focused in a few months. Linear ICs are still usually designed by a small team of 1 or 2 engineers. And, unlike the reliance of digital designers on cook-books, standard cells, and libraries of basic building blocks, linear ICs are usually designed from scratch and customized because the performance depends critically on the details of the transistor and resistor design and layout. This must come under the purview of one person. The point I want to make is, that the techniques necessary for successful digital design and for successful linear design may have some areas of overlap, but many areas where there is nothing in common.

Planning

Whether we do it explicitly or implicitly, every good linear design starts with a data sheet. If a circuit is worth doing, that is because eventually an engineer will pick up the data sheet and say, "I really have to have this circuit, it's exactly what I need". The circuit must include features and specifications and applications circuits which are excellent and will exert a magnetic pull on the prospective user, and these features must make a good total package when put on a data sheet. You don't have to write up the entire data sheet at the start of the project; it may be enough to just envision it, and write down the key features, but the designer must be aware that the data sheet will not only be the birth of the product, and the life of the product, but, if the data sheet is deficient, it can be the death of the product. So, the designer (with some help from his friends in the marketing department) must be sure that the circuit can lead to a data sheet that will have all the right features, and not any significant drawbacks. A single specification is not enough. You have to have a reasonable package of specs and features, and you have to have applications notes so the user will not be able to fail to apply it successfully. At every stage, the designer (or the project manager) must be aware of the nascent data sheet. Any decision that would impact the data sheet must be considered. For example, if a test engineer begs to drop a test because it is too difficult, and this would impact the data sheet seriously, the designer must resist this. [1]

Of course, the "silicon" must be able to do its job, but after that is accomplished, the data sheet must do its job. For example, the old UA709 amplifier had some good features, and some characteristics that were not so good. But the UA709's low noise was never mentioned on its data sheet. So, to this day, people ask me, "Where can I find a low-noise op-amp with lots of gain-bandwidth product?". And I have to explain that the UA709 really does have good low noise voltage, lower than all the 741s or the BIFETs. So, in some applications, the 23-year-old UA709 is still the best IC for the job. But the data sheet kept that a secret.

Design Reviews

When the design of a circuit is nearly complete, the designer will normally hold a Design Review for an audience of his peers. The complete design is presented, and all the engineers are invited to marvel at the good design, and to criticize the bad design. On a good day, nothing is left but the good stuff. But, we all know that occasional design flaws slip past the jury of peers. To help minimize this, at our company, we have appointed several Czars who are intended to be experts in one particular narrow field of design. I am, for example, the Czar of Bandgaps, and I am responsible for keeping track of all our expertise in the design of Band-gap reference circuits. Every new circuit, every old circuit, every good

circuit, every bad circuit, and every "fix" of a bad circuit is supposed to be filed with me. That way, I can make sure that only good circuits are copied. This theory does not work perfectly, but it works rather well, and the number of old errors that are repeated has decreased markedly, recently. Czars have also been appointed in other areas such as: Start-Up Circuits, Zener-Zap Trim Circuits, Proof-reading, and Computer Errors. [2]

After the nominal circuit design is correct, it must be laid out by a mask designer (though sometimes by the Design Engineer himself). Either way, when the layout is complete, it is necessary for another round of detailed criticism, peer review, and close checking. This we denote by the technical phrase, "Beer-check". This differs slightly from the Design Review, as there is opportunity for all the technicians, and mask-designers to review the layout, as well as the engineers. Now, as I mentioned earlier, there may not be as many bad or incorrect ways to lay out a digital circuit as there are bad ways to design it. But for linear circuits, there are a lot of bad things possible to do in a layout, even for a good engineer and a good mask-designer. The kind of flaws that are spotted at almost every beer-check range from the trivial to the disastrous. Fortunately, we often find that the most unlikely persons have the best knack of spotting errors, so we invite the most unlikely people as well as the senior persons. After you try this a few times, it is easy to conclude that the cost of the pizza and potables that are given out as rewards, is 2 or 3 orders of magnitude cheaper than letting the errors go un-detected. An even more important feature of this review, is that every level of worker learns that every person's input is valuable for catching errors, and that nobody is immune to errors. It is especially apparent that when a person has worked on a project for a long time, he becomes incapable of catching his own errors, so that is another reason to bring in outside help, simply to bring new eyes to the task.

Testability

In digital design, when the number of gates rises into the thousands, the importance of testability is well appreciated. So, gates that might fail and go un-detected because of a lack of accessability are given special consideration. Likewise, in linear circuits, the importance of being able to test every important function easily and quickly is recognized - especially in the case of mixed-mode ICs where many analog and digital functions are combined. [3] Also, the techniques for designing testable linear ICs are being expanded. This is true not only for production testing, but also for characterization and for troubleshooting. Every node must be accessible to the troubleshooter with his probes, or else that becomes an invitation for Murphy's Law to strike. [4] Small probe pads are best introduced, to facilitate probing, and vias must bring up nodes from the first layer of metal to the top layer of metal. Fortunately, new instruments are also helpful in aiding the access for probing, but these are tedious and slow.

Computer Simulation

The advantages of the (digital) computer for simulating a circuit's performance are well known and widely trumpeted. The disadvantages are most often swept under a rug or ignored. Many young engineers are slow to catch on that computers lie. Yes, they sometimes lie even when you type in the correct SPICE program. Consequently, wise designers learn to run "sanity checks" on their computer programs, and to blow a whistle if the computer gives a patently absurd answer. Other times, the computer itself blows the whistle and gives up - quits - fails to converge. In every case, the engineer must be prepared to evaluate the results and to decide when to dis-believe the computer.

In a recent case, an operational amplifier was designed with MOSFETs, and all the DC characteristics were simulated and the results were as expected. However, the high-speed response and dynamic stability tests gave no results, due to lack of convergence. Despite pleas, howls, and threats, the Computer Experts were not able to achieve any useful convergence. (Note, this may be partly because the MOSFET models are optimized for accuracy in digital circuits, where the voltages flash from maximum to minimum in less than a nanosecond, but the models often give poor results in linear circuits....) Finally I suggested an analog computer. The designer said that would be difficult as he did not have any suitable kit-parts or test patterns - no suitable samples of MOSFET. I pointed out, that's no problem; replace each MOSFET with a bipolar transistor with a big resistor in its emitter, to simulate the transconductance of the FET. He thought about that, and then objected that the bipolar transistor would have unrealistic frequency response. I replied that he should add in capacitors 1000 times bigger than the actual IC strays. Then the circuit would work at 1/1000 the speed of the real circuit, and the stray capacitances would have negligible effect. He thought about it. He convinced himself. He built the "Analog Computer". It confirmed his theories. He built the IC and it worked just like the Analog Computer. So, when the digital computer in the simulator gives answers that are useless, you are not necessarily stuck.

Recently we had a large non-linear circuit whose simulation gave absurd errors. Fortunately, we were able to snip away 63 of the 65 transistors and the SPICE continued to give absurd errors. By concentrating on the minimal circuit, we were able to convince the Computer Experts that they must

look in their realm for the error, and in only a couple weeks they were able to find 2 or 3 absurd errors. Then we went back and the complete circuit did begin to work correctly.

Breadboards

As an old-time analog-circuit crank, I find that breadboarding is a valuable part of circuit design. For one thing, it does provide an independent check on the results from digital simulation. For another, it can use "kit-part" transistors so if the transistor model is incorrect, the circuit may tell you a different answer than the SPICE, because the model is poorly related to the real world. But, just as I cautioned you that digital computers lie, well, so do Analog ones, and breadboards, also. They lie most easily in high-frequency response, where the stray capacitances of the breadboard cannot relate to the strays in the monolithic circuit. Still, a breadboard is valuable, I find, exactly because it is so touchy, so sensitive, so grouchy. These days, it will normally be foolish to depend only on the breadboard, but in many cases it's not much more foolish than to rely solely on the computer simulation.

Another good use for the breadboard, is to exercise the tester. Several times, I have been able to back up the breadboard to the tester, and confirm that the tester was working OK, even before we had working silicon. So, when you want to run a "sanity check" on the tester, the breadboard can be valuable and reassuring. In some cases, you can see the effect of a sweeping change in the breadboard, quicker than on the computer. So, in many cases, I recommend that you should not abandon the old techniques, if they help you to avoid the occasional hoaxes that the digital computers attempt to foist on you.

More Planning

The way to use PERT charts and Gantt charts to best manage the progress of a project is well documented in the art of program management, and I will not dwell on that. I do want to point out the advantages of using a big Check-list. I have been using Check-lists for a number of years to make sure I didn't forget anything, because on a complex project, there are so many things that could easily be forgotten or neglected.

For example, on a recent project, I started out one evening to write a list of "100 Items to do", and when I was done, I counted 155 items. On another project, I was collaborating with two junior engineers, and each of us wrote down a long list. I had on my list more than a dozen items they had not thought of, and they each had a dozen items I hadn't thought of. So, by combining lists, we came up with a much better list, which will also be useful on other future projects.

Priorities: In these lists, we try to indicate the relative priorities of various aspects of the project. That way, if the noise is the most important feature, and you find yourself compromising the noise to gain other advantages, at least you cannot say that you didn't realize the noise was important. One time I was given a new project, and as I wrote down the list of 100 Items, I tried to decide if it was most important to keep the die size small, or the quiescent power, or the tempco, or, what? (This was because the target specs and guidelines were rather imprecise.) Shortly, I decided that none of these had the highest priority, not even, getting the silicon out in the fastest possible time. Rather, the most important item was, to have a 100% probability of getting the circuit working on the first try. This circuit was to be part of a new library of cells, and the other cells also had to come out on the first try. Some of those cells had 8 components, and others had 18 components, and my circuit had 85 components, so it was not easy to build up the confidence that my circuit would be sure to work on the first try. So, I optimized the circuit in terms of getting each component precisely defined, with no mistakes. And, with a little help from my friends at the Beercheck, we did get it out on the first try, and it worked well.

Now, if it had not worked on the first try, what would be the next most important thing to optimize? In this case, I should optimize the ease of analyzing any errors in the first try, and the ease of fixing them and getting the second silicon to work. So, I also did that. And, even though the silicon did work on the first try, I was able to make a minor metal-mask tweak to make it work even better on the second try.

Why Metal Mask Tweaks? There are a couple advantages in being able to tweak the metal mask. If I have a circuit which works pretty well, I can usually predict how to improve the characteristic nicely by changing a resistor value. So, I could achieve this by changing the base mask. In theory, that is easy to do, but in practice, it's easy to botch the computation, and get resistors that don't do what we want. If, on the other hand, we build in a large group of small resistors in series with the main resistor, and short some of them out with metal links, we can prove that the resistor can be changed to the desired value by opening or shorting links. We can do that on some samples, and then we know that when we change the metal mask, the improvement will be exactly what we want. Further, we can do all sorts of evaluation on these parts, (which would otherwise have to wait until the base mask was changed).

Another reason to make the changes on a metal mask, is that we can more quickly get the changes executed. If we hold some wafers at base (before base mask) then we can slap on a new base mask and get out new wafers in 2 or 3 weeks, but if we hold them

at metal (before metal mask) we can get new wafers in 2 or 3 days.

"Hold At Metal" The other major reason for planning metal mask changes is related to major goof-ups. Many times a new run comes out, and absolutely nothing works. A little study shows an error in the metal mask. We can go back and start new wafers, and put on a new metal mask, at a cost of a few weeks and many thousands of dollars. If instead, we hold 9/10 of that run at metal (before metal mask) and bring out 1/10, we can learn from looking at those few wafers that the metal mask was botched. Then we can expedite the new metal mask, and slap it on the left-over wafers, and cut down on wasted money and time. I have not declared myself the Czar of "Hold at Metal", but I am one of its strongest proponents, and I have convinced everybody I work with that it saves money, as well as time which is more priceless than mere money. Now, it is true that having metal-mask options available on resistors does waste a little die area, but I've seen this cost paid back 10 times over. And the cost advantage of holding wafers at metal could be computed at 100 times or more.

"Bust-proofing"

Ideally, all ICs, all circuits should survive the outputs and the inputs shorted to either supply, and large ElectroStatic Discharges (ESD) to any terminal. But, to be realistic, there are some circuits whose performance would be compromised by any attempt to make the circuit un-bustable. For example, the old UA709 would be destroyed in a few seconds if its output was shorted to either supply, and even faster than that if one input was shorted to a supply. The demise of the UA709 was easy to predict when the LM101 and UA741 came along, as they could survive all these conditions. However, recent requirements that all IC terminals must withstand ESD transients larger than 2000 volts (standard test procedure, applied through 1.5k in series with 100 picofarads) showed that the LM101 had weaknesses when the inputs were pulled toward -80 volts. We were able to add a couple small transistors to clamp the negative-going transients and bring the LM101A's toughness up to an acceptable level, and the user cannot even see the difference. Of course, customers still call us up and ask how many transistors the LM101A has, because they believe MIL-HDBK 217B, which asserts that the circuit will be more reliable if it has fewer transistors. We know better than that. The LM101A has become MORE reliable for the last couple years, not LESS reliable, because it has had 2 transistors added. In general, adding transistors that make the circuit more bust-proof, more forgiving, are good investments in customer satisfaction. In the design of a "high-volume" linear circuit, there are many cases where this is proven wise, to add circuitry in the interest of "foolproofing". But, of course, we don't call it that.

Testing Strategies

Sometimes a test engineer tells me, "Here's a test that almost never has a failure, so I want to delete it." Oh, what is that test? It's a quiescent power drain test at a 6-volt supply. But, that's odd, we know there are a lot of parts that are dead and would have to fail a quiescent current test. It turns out that the 6-volt quiescent current test is run after the 33-volt quiescent current test. Consequently, that test lumps together the dead parts and the ones that break down when the supply is increased from 6 to 33 volts. By swapping the two tests, we can learn quite a bit from these tests. By putting the tests in a foolish sequence, you fail to get your money's worth from these tests.

Recently a product engineer asked me why we were getting such poor yields at final test. I checked into it and discovered that the wafer-sort test had much wider limits than the final test. By being too busy to set the limits in a rational fashion, we were condemned to spend the costs of packaging up bad dice and throwing them away. Normally, it's rational to test and throw away the bad dice before assembling them.

Other times when a test engineer complains that a batch is having a terrible yield on one test. I ask him, "Are they failing any other tests, too?" Often, they reply, " I don't know, after this test, the parts failed and weren't tested any further". I have to explain why they must test these parts in "ALL" Mode, rather than "PASS" mode, so we can learn about the results of the other tests. Usually on a well-designed part, you can learn more from studying bad parts than good ones. I usually spend a minumum amount of time studying tests that give good results, but a lot of time studying parts that fail tests. Sometimes I even have a hundred "inked" dice bonded up. You can learn more from studying the sinners than the saints.

Summary

The world is a complicated (and non-linear) place.

References

[1] R. A. Pease, "How To Get The Right Information From A Datasheet", EE Times, April 29, 1985, pg 31.

[2] R. A. Pease, "Preside Over Power Components With Design Expertise", EDN, Oct. 12, 1989 p. 134

[3] P. Fasang, D. Mullins, and T. Wong, "Design For Testability For Mixed Analog/Digital ASICS" presented at IEEE Custom Integrated Circuits Conference, Rochester NY, pp 16.5.1-16.5.4, May 16-19, 1988

[4] R. A. Pease, "Troubleshooting is More Effective With The Right Philosophy", EDN, Jan 5 1989 pp. 147-156.