ISSCC 2011 Tutorial Transcription
Layout - The Other Half of Nanometer Analog Design
Instructor: Jed Hurwitz

---

1. Introduction

Thank you. A couple of quick notes, this first one is about mobile phones, so if you could, please put them on silent, I'd appreciate that, I've just done mine myself.

Next one is there won't normally be any questions during the tutorial but I'll be around afterwards. So if you want to corner me, you're welcome. I don't know whether I'm going to live to regret saying that, so better get started.

So first these are my acknowledgements. I'm not an academic. I don't teach for my living. I do do a lot of PowerPoint engineering in terms of technical marketing, but to pull together a presentation, for this I had to go to some of my friends in different companies and ask for bits and pieces of material. Hopefully you'll find them all useful and

2. Intro

I hope I haven't missed anyone out there

3. Acknowledgements

who helped me pull this together.

4. Tutorial Aim

So firstly on the tutorial aims, it's really meant to look at layout going forward into nanometer technologies. There's a lot of stuff around about analog layout in general so this tutorial is going to focus on some of the issues and the effects and some of the challenges when you go into nanometer technologies. And I'll try to introduce some of the issues associated with metallization, how you could maybe approach layout to shorten design cycles and make them more robust. And as with anything, the best way to learn unfortunately is through problems or mistakes so I'm going to highlight some real design examples where hopefully it will teach you through a kind of analysis of a real problem, what type of thing you need to be aware of, of course there's a good chance it won't be applicable directly to any of your own circuits. What I want to teach you here and give you some tools is to be aware, to be paranoid, is a very useful thing, being an analog designer to be paranoid. What I'm not going to get into is structures to make various passives or components, particularly RF, I'm not going to really get into substrate noise and coupling between digital or analog although I'll mention some things about digital within analog. I'm not going to mention stuff that will happen below 40 nm where things, the physical effects around the transistors are going to start changing again because that's a whole new ballgame to start learning.

5. Disclaimer 1

So I'll start with a disclaimer. Every circuit is different, every technology is different, even the same geometry in a different foundry can be quite different, every requirement spec is different, every CAD flow is different, and every designer is different in terms of their approach to designing. So it's your problem guys,

you can't take the stuff you're going to get here and "Jed said this, that's what I'm going to do, it didn't work, it's Jed's fault", I'm not going to take that, it's your problem, and that's why I hope you will learn from this.

6. Background

7. Why is layout important to a designer?

It's a very competitive market place we've got in terms of integrating functions, less and less can you afford to put in a function purely to do the job, you have to be at least competitive with other companies, or with the stand alone solutions, so if you're going to embed analog, it's got to be relatively high performance, now power speed and precision is not all about the design, for me a design is the topology and then you choose the topology for a function and then you tweak the components within the topology to maximize on function, but for a given topology there is some theoretical limit. And the closer you go to the highest performance aspect of that limit, the more you'll be at the mercy of some layout effect causing you a problem. The layout does cause problems for headroom, systematic mismatch is what most people associate with layout, and crosstalk, and even gain bandwidth layout can effect.

Next point is risk. Your design cycle is very long and you spend an awful lot of your time working on why it's going to work. There will be hundreds of things that you need to get right for it to work. The unfortunate thing is you only need to miss one thing for it to be a problem. It's a kind of unfair waiting, but that's the reality. One of the challenges is not everything is modeled, and not everything is extracted, and not everything is simulatable. So we use the tools because they help us do more and more complex things but we have to be aware, they have their limitations and we are designers who have to be able to put something else over to be able to look at the failure modes and understand why things fail, that is much more important to why they work.

Design time, there is more and more pressure to get the designs out but if you actually look at the plot of all your projects you've ever worked on, the main reason for the long project is not how quickly you go through the path, but how many times you have to go around the path to get it to work.

So basically, the motto here has been right first time, it saves you in the end more times than anything else you could possibly do.

And then area, analog is not shrinking as proportionally as you'd expect, particularly things that interface the real world, whether it's PAs, line drivers, etc. So, there's lots of pressure to squeeze them down in size. They can be equally problematic to try and blow things and make them bigger. We'll talk a little bit about those effects.

8. CMOS Analog Design Relies on Matching

I'm going to start with some very fundamental things. Firstly a very crude statement which is CMOS analog design relies on matching. And the very first part of the matching, most people associate things with is input offsets for an amplifier, or a comparator, but error here could also affect things like common mode, errors at the input stage could also affect your common mode rejection ratio and your power supply rejection ratio. So if you simulate your perfect amplifier and you don't put in any error, you would probably have a really high common mode rejection or power supply rejection ratio. It will degrade in some way when you include some of the offset affects and not everyone appreciates that.

Two aspects of offsets, the pure millivolts, and then there is the impact on other parameters. Likewise, an awful lot of matching studies go on about precision, whether it's for a converter, what kind of resolution can you achieve, the effect it will have on your INL and your DNL, but it will also have, mismatch will also have

an effect on other parameters like SFDR. Maybe not within the converter itself but its output stage and other effects by bringing in other harmonics, it happened to be cancelled when you had a perfectly matched system.

For me, one of the most important areas that relies on matching nowadays is biasing. So it's getting trickier as we go into nanometer, the headrooms available are not as high; we've got to make it work over process-voltage-temperature (PVT).

9. CMOS Analog Design Relies on Matching

Your matching effects your operating headroom, effects your quiescent current control, more importantly as we go forward into very thin oxide devices, it effects your safe operating area (SOA) which can maybe effect the lifetime of your product. So biasing matching use to make up biasing is becoming more and more important.

10. Mismatch is NOT ONLY Caused by Microscopic Device Fluctuations

So here is a corny slide, basically two matches representing matching.

Most studies or most theories around matching is around microscopic fluctuations in the edges, the materials effecting the matching and so you have an area term. For the reality is your devices won't necessarily match if they're not used equally. That means you might have a difference in the biasing you use for them, you might have a difference in the local temperature, you might have some electrical disturbance from a neighboring circuitry that causes a difference. There, the reality is that devices are not as equal as you expected, so maybe you've put them in different orientations, maybe you have a different environment around them that's affecting them, maybe you've got different metallization over them. Now the first one is pretty much a layout and design challenge. To some degree, you've got no excuses, you're making the decisions around that.

The one on the right, you have to be aware that some the effects are characterization/modeling related. So if you're not aware of them, if you've not been able to extract them, maybe you've got a problem you don't know about.

11. What Do We Want to Know About Mismatch?

So, there's lots and lots of studies out there on matching. I said the most important one is the A-factor, the area factor. Pretty much it all matches theory, larger things match better than smaller things. It goes to the other extreme where in the end, you've had to do dedicated mismatch investigations because you've got some unusual failure in a product and you need to find out why it's not working.

As we go down this, unfortunately there's less and less literature on matching and the effect of the physical side on what you get.

12. What Affects Matching

So again, going back to my corny slides, what affects matching, the orientation? So ideally you should have the current flow in the same direction, you could have them in different directions, or orthogonal, they could be diangle. This is the one that most people associate with that gives you best matching.

Now to achieve that when you have multiple components, quite often you'll use multiple fingers, you might have them facing other ways, you might lay them out in a common centroid fashion in a line array or as a quad match fashion like that.

Another factor that affects matching is distance and whether things are close or whether they're far apart. I'd say as we go down into smaller geometries, there is less and less difference between things being close and things being far apart. It doesn't matter as much, the process control is getting so good or you could say it's getting so bad, that it's more equal as you separate things.

And then you have the affect of overlaying layers, whether it's layers that go over your transistors or adjacent layers next to transistors, they will have some kind of effect.

13. Layout Can Affect Power Efficiency

Layout can also affect power efficiency, so crudely speaking Gain-Bandwidth for a given current is dominated by what kind of Gm can you achieve and what kind of parasitic capacitances do you have in achieving that.

Here's an example of three different layouts, they've all got the same W/L, but actually because of stress affects, because of capacitor affects, they all have different Gm's and Cs's and Cd's, and some of them will be better than others for those affects.

And this can make a difference in your circuit performance of a few percent, maybe even up to 10 or 15 percent; you can squeeze out of actually using the fact that layout impacts performance to your advantage.

14. What a Designer Should Really Care About

So, coming back to that first side of it is, what should a designer really care about? Now, I could talk about the first few things are kind of obvious but more and more important is how long will it take to get the design ready to tapeout? Will the design work when it comes back? And will the design have issues as it goes through qualification and production.

15. Classic Layout Techniques

So I'm going to start with some classical layout techniques.

16. Classical Analog Design Flow

So a typical analog design flow has a design phase, and hopefully somebody has given you some specifications up front of what you want, maybe you've developed them yourself. You look at the specifications and you choose the topology and around the topology you try and optimize the W/L's, and the resistors, and capacitors and bias currents and you enter into simulation phase and you go around this loop until you're happy with it, and then at the end of the design phase, you kind of pass it over, sometimes over a fence. Some companies are very regulated, designers do this and layout people do that. But in other companies, it's much more integrated so there's a difference. You pass it over to somebody who enters the layout phase and they really take it through to a kind of DRC/LVS clean and then you go into a validation phase where you extract and simulate and you then see if the performance is met, and if not, you go back around the whole loop. And hopefully that should be familiar to you as a basic design flow

17. Classical Analog Layout Techniques

So the classical analog layout techniques to achieve matching are common-centroid. This is an input pair out of Willy Sansen's book, it's actually a single poly, single metal type design. It's very old, it's been going on

for a long time. It kind of works, but is it always needed is one of the things we should be asking ourselves. Do we always need to go down a common-centroid?

18. Classical Analog Layout Techniques

because as we go forward, and you look at it in terms of a system level, this would be for say a 4-bit DAC, including the LSB of the segmented DAC in the middle and you're trying to put all the other LSB's around so they have an equal grading going across, the question is, does that introduce more problems than it's solving? It's very easy to take in an input which says "I am going to common-centroid everything", but it can become a bit of a nightmare. We'll cover that a little bit in some examples going forward. In the end, you have to decide, is it bringing you anything or is it actually introducing more problems?

19. The Folded Cascode Example (1)

So, I'm going to use an example of a topology, a fully cascode amplifier, very common. You should be able to look at schematics and identify what you might want to be optimizing in your layout. So you can go to the input pair and of course we want to match the two input pairs, you maybe want to match also the Pos/Neg output to improve your power supply rejection, your common mode rejection, your balancing of the output and also any output offset that becomes input referred.

You're also matching the biasing circuitry to the rest of the amp and then within a given transistors, you'd really like to optimize the Gm of the Input pair and you'd like to minimize the capacitance of what's your secondary pole and what's your primary pole. So you should be able to look at any schematic and identify what things might translate from schematic to layout.

20. Matching a Topology to a Layout

If I take this example and go forward, you can say these devices here are your P-MIRRORS, these devices here are your P cascodes, these are your N cascodes, these are your N-MIRRORS, these are your N input stage and I'm going to build it up like this in my layout.

And you notice here, I have chosen to put the input, not in the same place where it is in the schematic. And that's because there's several things you need to consider, it's not just trying to map your schematic to your layout, it's also, what are the questions to be solved here? The output node is between the P cascodes and the N cascodes. And from a capacitance point of view, it's not very desirable to have your input close to your output where you've got a high gain in between so any parasitic capacitance would cause you some kind of problem. Some people put the N Input underneath here. There's various different layout strategies but in the end, I didn't map it directly from the schematic to the layout just because I drew it that way in the schematic. I made a conscious decision to make some alterations.

21. What Happens in Nanometer CMOS to Make the Layout More Challenging

So hopefully, not too much of that should be that new to you. There are some books and documentations about what I call the classical layout stuff.

Why is it becoming more challenging as we're going down nanometer CMOS? Part of it is that there's more pronounced physical effects. In particular STI stress, well proximity, poly proximity, metal overlap, also metallization is becoming more tricky so metal and contacts is becoming more resistive, there's lower electro-migration and current handling capabilities. And then there's a whole set of design for manufacturing flows that you need to accommodate in your layout that actually alter what you thought you were doing to what

you need to do to complete the layout.

We've also got low voltage headroom, and that means that we've got a smaller SOA window, and there's more use of digital assisted analog. Generally that means that there are more signals, and that there are more wells, and more supplies to handle. The question is, how do they interact?

22. Things about Nanometer You Need to Know

23. Shallow Trench Isolation – Stress Effects (1)

So I'm going to talk about some of the effects. Shallow Trench Isolation, it's actually been around, probably, mainly for around 0.25µm. Certainly some people modeled around then. It's gradually taking on a larger effect as you go down into 90, 65 and 40 nm. It's quite well modeled, that's a good side of it. The basis of it is that the edge of the active, there's some stresses here that is a function of the fact that you've got a change in the physical side and this has an effect in the strain of the silicon as you go along and therefore the transistors have a different characteristic as you go along. It's incorporated into BSIM4 SPICE and beyond and is basically three parameters SA, SB, SC that model it.

24. Shallow Trench Isolation – Stress Effects (2)

There's a very simple equation here as to the base level of what stress you're getting.

25. Shallow Trench Isolation – Stress Effects (3)

And the real thing is to realize that it can have quite a huge effect, so here's a little study that somebody did where you had 15 percent of Gm with three different positions away from where the edges of the active. The good news is that it can be extracted, but the bad news is, don't leave it to the end of the design. You're going to have some assumptions into your schematic about what SA, SB, and SC you're going to have for each choice of transistor inside of your topology. That's going to be there primarily from your choice of number of fingers and number of units of your transistor and the Pcell will self generate that. When you come to your layout, you need to make sure some of these things are matched. Otherwise you're going to have to go back in that loop around the layout when you extract it, is going to be different by up to 10 or 15 percent which is quite a large amount.

As I mentioned before, you can often use this to optimize your transistors. So naturally, say like an input pair or for certain switches, you might want to maximize the Gm, this might want you to use the effective structures that have a very short overlap of the active to the poly.

26. Can Dummies Solve – Shallow Trench Isolation Stress Effects?

So, a lot of people say, can't you just put a dummy in to solve your problem? Just put a dummy on the end of the transistors, it will solve it. This is just a simple test structure I laid out and extracted, just put a whole row of transistors down and looked at what, biased them with constant gate, and looked at the effective error in current as I moved along in distance. And you can see that to get to 0 effect, you had to be 4 µm away, and that's quite a big transistor for a dummy nowadays where a lot of your transistors, even for analog, are sub microns or 0.1 µm. So just going around saying, I'll put a dummy on, is not the right way unless you can justify it. I'm not going to say no to dummies, but I'm going to say that you should always be aware to what your technology does.

27. Well Proximity Effects (WPE) (1)

This is a similar effect to Shallow Trench Isolation stress which is well proximity effects, really it's a kind of well overlap, so it also covers well overlap, but the well can affect the dopant gradient, and that alters the transistor characteristics.

28. Well Proximity Effects (WPE) (2)

Again, can a dummy solve it? The answer is kind of yes, but again you have to be aware it can be a few microns before you minimize the error associated with it.

So if you'll actually look at design rules, say from some of the major foundries you'll quite often get digital design rules with a certain well overlap, you'll get analog design rules with another overlap, you get digital design rules from manufacturing with another overlap and analog designs for manufacturing with yet another overlap. And each of them, the overlap size will start to increase, and that can make quite an effect on the density of your circuits, particularly if you want to interleave PMOS and NMOS and different types of transistors in your layout, you'll find gradually your spacing is increasing, dominated by this Well term.

And so for myself, I'm not recommending this, but it's opened up a question which was, do you really need those design rules if you're fundamentally using matching? Can you go to the minimum and can you accept that it will change the transistor as long as I know what it is, and as long as it is the same as my neighbor, is it acceptable? In the end we accept it on a compromise, we did something in between. But the only way to prove it because there's no documentation, is to put some test structures down on it.

So we actually have that inside a product in a fab at the moment, so I'll find out in a few months time whether I can break those rules. At the moment, I can't tell you to break those rules. It's your own responsibility, but any rule is there for a reason, you need to understand that reason and you need to think about what its impact is, and is it simulatable, is it extractable, is it measurable if you want to use it, if you want to break that rule. So, it's just an example of rules are there for good reasons. They're there to make sure things are safe, but sometimes you have to decide whether you can look at them.

29. Poly Proximity Effects

So poly proximity effects, there are some stress effects, poly nearby, there's process effects which come from the edge and proximity if you have a large area of poly nearby, you can have some influence on its neighbors, and there's some optical effects. A lot of these are compensated for mask-prep. The only issue about them being compensated for mask-prep, is they're outside your control, so the safest way to know that things are going to be not affected by this, is to put down, make sure you've got even structures going along. If anything, you want to match, and in this case, you need to put down some kind of dummies around so anything that is going to be automatically generated, the differences are going to happen on the components that you don't care. Because they're actually outside your control, they're being done to maximize digital yield and digital leakage and the very good algorithms associated with it is the optical correction for the film and you've got to be happy that they're not going to influence anything that's your sensitive circuitry.

30. Metal Over Transistor

So I'm going to raise one here which is metal over transistor. Some of you will be aware that some foundries or some processes say, you're not allowed to put metal one over the transistor, some of them are, they don't tell you anymore than that.

Now what's the origin of this? I actually found this out for imaging where we push the requirements on leakage and on transistor 1/f noise a lot and we have arrays where of course in imaging, you can't put metal over the arrays but on the peripheral, you're going to have some metal. We notice in effect, which are the

pixels that are near the peripheral were different to the pixels in the middle of the array. And the reason they were different wasn't because of any poly or active regions, is because the metal density near the edge was very different. And we found in the end, it was because of hydrogen annealing so when the process gets finished, the very last step, they go through a 400˚C typical anneal and this releases hydrogen from the passivation which is part of the annealing process and this has an effect on the kind of defects within the substrate.

Now it's all very interesting in imaging, I doubt whether most of you people here are related to imaging, but later on I found that this effect is there also in and around analog and it's not just a metal 1, metal 2 effect, it goes all the way through the metals stack.

I've searched in the literature and there isn't much stuff about the metal effect on the end-of-the-line transistor substrate, there is some stuff from the bipolar days where it had a stress effect, where a neighboring metal could alter the stress in the substrate and all through the B terminal of your bipolar, but again it's an effect that's not extracted and not modeled and the basic thing if you've got an array of transistors like in a DAC you need to match. You need to make sure the metallization around each of them is the same because it can have an effect.

31. Metal Resistance, Time Constant and Electro Migration

So I mentioned about metal resistance, so firstly I would just like to make a bold statement that you can't treat metal like metal, treat it as a resistor. Stop calling it metal as if it's got some magic property, which means you don't need to worry about it. And that's because it affects more and more the IR drop, the RC delays and the cross-coupling interactions. If you look at the cross sections of metal nowadays, there's more side plate, more in the sides than the surface area of them.

And just giving an example, even just within nm, going from 90 to 40 nm is almost like an order of magnitude delta in RC, for the same minimum track. When you, for example, are laying out some digital control lines and passing some analog signals, you tend to do things using the minimum wires you can get away with for the current density.

For example, the contact resistance is going up 2½ fold, the resistance for the minimum track is going up 10 fold, the via distance has gone up 10 fold, and so forth. So what appear to work in 90, may not work in 40 unless you are aware of the fact that you had an order of magnitude to change that. Maybe what you have to do, is increase the size of metals you're using, use more contacts, more vias around the place.

32. Poly/Active/Metal Density & Other Litho Effects

So, there are other effects, basically you'll always notice as you go down in technology, there is more and more density design rules. And this is to try and make sure that the process delivers something uniformly. Basically, the layout on the design is now helping the chemistry of the processing. So it is checked on poly and active, as well as metal. It's checked on different size windows. For example, 100 x 100 μm, or 10 x 10 μm, they might have slightly different density limits. You can get a result where you start, you do your design, you do your layout, you run it and it's clean, and you place it 20 μm to the right and you run it again, and now it's no longer clean. Because in the end, the sum algorithm is running a kernel, that's running over your design.

So, there are auto fill tools available to help you, but why I'm raising this as an issue, is that you need to make sure that your block doesn't change when you place it inside the macro. So you've done all this lovely work around your amplifier and the switch cap circuitry around it, and then you're going to place it within side your ADC, and all of a sudden, it's not going to be design rule clean, and then you're going to have to go and

change it again.

The other one is, what's the effect if you leave it to the auto tools, so you do your sub block and then you place it in your macro and then you press your button. I've already talked about how metal will change the characteristic of your component, so are you sure that the auto fill is going to not cause you another problem. It might make a design rule pass, but is it going to cause you a matching problem?

33. What Happens to the Metal Fill, Do You Get Extra Capacitance?

So for example, another side of metal fill is, you might have purposely separated the 2 tracks from the top plate and the bottom plate, the bottom metal, in order to get you the least capacitance where they were. And then you press, you do the metal fill, you introduce lots and lots of different pieces of metal in between that are floating. And this leads to a dilemma of what is the actual capacitance that you've now got. Are you introducing some cross talk path from the metal you've inserted, that you're deliberately trying to keep away.

So the best thing to do around metal or metal fill, if you have sensitive circuits, like switch cap circuits, or feedback loops involving capacitance, is to make sure you put in the metal fill yourself, design rule satisfies, then to make sure you know where the metal is when you've got lines that you want to reduce the capacitance. And make sure that it's still clean from the design rule point of view because you can't rely that the auto fill tools won't introduce you some problems later on.

34. Summary – How are These Effects Covered

So here's a kind of summary of how the effects are covered, so stress effects, they're pretty well modeled, they're pretty well extracted, same with well proximity. Poly proximity is not really modeled that well, it's not really extracted that well. Metal over transistor again has that problem, it's not modeled, it's not extracted. I put down here, one that I was actually surprised about which is transistor orientation isn't actually modeled at the moment, I believe on the 28nm and beyond, it will be. But you can't actually say, if you were to put down two transistors, one at 90° and one at 0° in your layout, both will be extracted as the same transistor type. And they would both match perfectly. I think most would feel that they would probably never match. But the reality is, the tools don't extract it and they don't model it. Sorry, I have put in here, you can now, there are some features in some of the extractions decks that as least let you, tell you that transistors are in different orientation but whether it would be validated, I'm not sure, it depends on your deck.

So track and contact resistance, they can be modeled but apparently not easily, because when you start to add in multiple bits of contacts and multiple bits of vias and multiple tracks to get around the fact of higher resistance, you end up with lots of nodes and the problem with lots of nodes is that your simulation starts to run a lot slower. So you have to do some form of reduction in your extraction to run some reasonable simulation with reasonable accuracy. I'll talk about that later but that causes a problem. So these things go hand in hand, the extraction and the modeling and seems to be the simulation is a challenge for both contact resistances and track capacitance with metal fill because you're starting to get exponential problems of number of components and simulators don't like that.

35. Digitally Assisted Analog

So, digitally assisted analog, say you start to look at this conference, the last few year's conferences, you'll see there's more and more, there's a change in design style, there's more digital related to an analog block than ever before. For example, trim bits on the offset biasing, this is like a little comparator that was part of a flash converter. We use minimum sized devices here to reduce the input capacitance of the flash. That meant that the offset for the comparator was too large for the resolution of the ADC we wanted, so we just

very simply put in a little DAC per our comparator, and we put in a little loop that happens in parallel of the converter, or each comparator goes and trims itself and that meant that I can have a nice low power comparator within a flash with a low offset and a low input capacitance because I was able to do that. But now I've got one DAC per comparator and I've got so many bits, and I've got so many other things that are going on, so the problem is, you've got more and more signal lines and that's not just the case within one block, it's also at the full analog front end. I'll show you an analog front end of a power line module later on but, we have something like 700 control bits going back and forth to this analog front end, 700 wires from the digital.

There's only a 12-bit DAC and there's a 10 bit DAC and a 10-bit ADC and there's an 8-bit ADC and there's another 9-bit DAC in terms of actual data signals, all the rest are bits and pieces of bits to manage power-down sequences, biasing options, etc. That's a lot of wires to be handled by hand. So it causes a problem.

Do you use digital or analog tools for including the digital. And the reason this is a big dilemma is, do you let an analog designer check the timing or do you let a digital designer into the physical side layout? Or worst of all, is some kind of combination of it. The other one is the most common issues here for analog designers and digital designers is this concept of a static control signal. A lot of the things may won't be that way. It's normally quasi static, you do in one phase of the utilization, you don't maybe care about the timing. But one of its issues, is that just because it's static, doesn't mean it's not a sensitive node in terms of cross talk or other facts like that. It's very easy, particularly analog front ends are disproportionally large to a normal digital circuitry with an equivalent number of gates. So if you have 10,000 gates of digital within a digital 40 nm device, there'll be 0.0001 mm^2, you could have 10,000 gates associated with your AFE, but the AFE might, those 10,000 gates might be spread over 5 mm^2, so it's quite a different thing, putting digital inside analog blocks, and putting digital inside digital blocks.

36. Sub Block Layout Strategies

So, I'm going to enter into some sub block strategies, these are more analog related sub block strategies, as to how to get more exact designs.

37. Transistor/Sub-Block Layout Style

Within a sub block, you might have, generally I like to think you should always split your transistors into multiples of at least two, part of this is that generally, for a most efficient area, most efficient capacitance, you should really layout transistors back to back. When you share a drain or share a source, because generally one of them is more important capacitively than the other. But also it can be important if you've got for example, like a pipelined converter, where you want to do scaling per stage, the first stage of the pipeline is your most accurate as you go through the pipeline, you can reduce the requirement on gain bandwidth for settling to achieve the overall accuracy.

Rather than designing five different amplifiers for it, it is often better to design one amplifier knowing that you want to modularize it in a binary fashion with 5° of scaling that's going to happen, because it will save you a lot of work later on in terms of validating each of the blocks.

And the other extreme of this, is sometimes I look at people's designs and they've gone down and chosen the unit transistor to be the minimum, 50 nm wide and 40 nm long, my input pair is 57,653 of these transistors and you go, "get real, find out what the best modularity is", and it's never hundreds and thousands of one small unit, and it's never one unit of something very large. And it should be an intelligent decision you make about that.

So basically, the main thing to consider is using single units of multiple transistors in your strategies. You might say my single unit is going to be 8 blocks of transistors that are 5 μm wide and 50 nm long and they're going to be interleaved in this way and that's my new unit. And anytime I want to double the current, I'm going to put 2 down of that unit. It's a little bit like making a macro transistor.

I'm going to talk about the kind of strategies regarding transistor, versus top level symmetry as you analyze, say an op-amp in what you want to do.

Basically, a transistor level symmetry is where you group together associated transistors like I showed you earlier in the classical technique. And top level symmetry is more where you lay out things as you layout one half and then you mirror it. And this is sort of two different approaches.

38. Arrays of Transistors vs. Array of M Units

In terms of arrays of single transistors vs. arrays of M units of transistors, a lot of people feel that there must be a very large area of penalty. Because the overhead of source and drain overlap is quite low, the actual overhead is not so large in area. You are more likely to match some of the strain effect by going down the array of M units and putting a dummy on the end to match it. You put one dummy on the end here, but it's unlikely that that's been sufficient. And the altering area is not so large.

39. Comparison of Symmetry at Transistor vs. Symmetry at Top level for a Folded Cascode Amplifier

40. The Amplifier (Schematic)

So, I'm going to go back to the folded cascode again. I put the biasing on the side, should be relatively familiar to you.

41. The Amplifier (Starting Components) Symmetry at Transistor Level (1)

And now I'm just going to say, well if you take that schematic, you can see that I've got an input pair, I've got my P-MIRROR and my P-CASC, you've got certain ones associated with the positive leg in blue and certain ones associated with a negative leg in cream. And then you've got these biasing ones on the side.

42. Symmetry at Transistor Level (2)

And if you are going to do symmetry of the transistor level, you could imagine you could do some kind of inter-digitation or common-centroid around your POS and NEG of the P-MIRROR and you put in the middle of it, your biasing transistors as they stay with it so that they're equal for both sides.

43. Symmetry at Transistor Level (3)

And you can work your way through and do that again for the N-biasing and N-MIRROR.

44. Metal Matched Cross Quad

For the input pair, I put this in here just to show you that particularly with metal densities and contacts and trying to get metal matching, it's quite common to put metal over the transistors, you just need to make sure that it's just totally symmetric for everything. Often that will lead to the most compact structure.

45. Symmetry at Transistor Final Layout

So in the end, you end up with something a little bit like a space invader. All the transistors are interlinked but you have good matching but you have poor parasitics because if you remember within for example, the P-MIRRORS, the transistors are spread across the whole width and there has to be some metallization that covers it to connect things together. So you're really trading off resilience to spatial gradients vs. adding in more IR, more capacitance type of issues.

46. The Amplifier (Starting Components) Symmetry at Transistor Level (1)

So the other approach is to take the same circuitry and look at it from a top level point of view and say that this input here and this stage here, like a pseudo single-ended structure and if you do that as one structure laid out together, and then copy it for the other side, you can make a kind of pseudo differential

47. Symmetry at Top-Level (2)

which is what I do here and you can put the biasing in the middle.

48. Symmetry at Top-Level Final Layout

And then you end up with something like this, so, it's pretty good for parasitic around the structure. It's got slightly worse gradient matching, but you could actually do something with a kind of mix and match. You could imagine that input pair across the middle there, you kept as common-centroid as a transistor interlayered and that dominates generally, your input pair, dominates for example your offset of your amplifier. You don't have to do one strategy or the other, but be aware that they have different pros and cons.

49. Which to Use?

So, here's a very crude thing, say in all the geometries, I found myself doing more of this type of layout, and in more advanced geometries, I find myself doing more of that type of layout. It's a generalization, but I think when you look at it, you'll see that the one on the right is more efficient with the issues of metallization and capacitance that you have.

50. Include 'Sundries' at Block Level Before Starting Layout

So, one other thing that I'm going to talk about in block layout strategy is, I don't know if people are familiar, I call it sundries, it's not associate with any sub block. There's normally a whole set of other things that you quite often leave to the end like decoupling, power down, trim, things you might want to put in for test, and debug signals or get out of jail functions, putting down some spares here or there, or bringing certain signals to the top metal so that you could look at them at FIB or change them at top metal because top metal mass are cheaper for doing any alterations or trimming.

51. Example: Getting Ready for Layout

Now, if you don't put it in, this is the same schematic of the folded cascode and you can see that you've suddenly added lots of components. So if you do all that good work of doing your floor plan with just the main components of the amplifier, and you haven't considered where all these other components are going to add, you've done that kind, I'll add them later, I'll find some hole and I'll put them in and if it's not in the right place, I'll root another track across to do it, you'll start to introduce problems that when you've got a nice stable design, you've done your extraction, it's working and now you add in all these other bits and

pieces, you actually need for most products. And then you extract it,

You suddenly find there's a difference. So, you need to include these sundries as early as possible in your design planning and your design flow. Because every time you go around that loop, it's painful.

52. Block Floorplan

So when you come to look at putting this sub block inside a bigger block, you've got to control the signal flow and what you do across the block, where do the signals arrive, where do they leave, it's very easy to make a sub block that is easy and will route beautifully but in the context of the next layer up, it actually is a disaster. You're passing the problem off to somebody else and that won't make very many friends.

Power supplies are often left as a global strategy to do afterwards. You say that I'll use the top two level metals for my main power supplies as I've talked about metal above transistors can have an impact, so if you've done a sub block and you want to use it within a macro block, and you said my VDD/ground is going to cross it, how are you going to handle that? Likewise for the metal densities.

Wells, as I mentioned before, they dominate a lot of your layout area, so again the organization, which wells do you associate with other wells? What well spacing and overlap do you want to use because probably, it's your largest design rule for many things.

And then finally a small point, many people will do things with deep Nwell, but be aware that not all designs are compatible with deep Nwell. So people have done some mix and match. It's very difficult at a late date if you're taking in a block from someone else, you're assuming it was going to be compatible. If it isn't, it becomes very problematic to find a way to cut out the deep Nwell around it, to take that block inside.

53. Building Up a Macro-Block Example Pipelined ADC

So, I'm going to give you an example of building up a macro-block around, say, a pipelined ADC, so you've taken the op-amp that you had before, around it you put the capacitors for doing the ×2 and reference retract and you've got the switches, and you could either do a structure, like these are top and bottom around the amplifier, or like these are before the amplifier.

What will effect things, it effects the compactness of the given stage layout and it effects the critical signal couplings back and forth.

54. Example (1) 12bit Pipeline ADC Floorplan

So, an example is to quite a common layout strategy for a pipeline converter where you've got scaling per stage, you've made a structure which is very similar as you go through your scaling. And you've routed it back on itself and some of the advantages is, it's quite a simple layout for each stage, also the signal flow is relatively clean between each stage until you get to the edge. And you've got to do this wrap around the corner. The other one is, this is your low-res flash at the end of your pipelined converter. It's very close to your input. You've also got some issues about how do you power all these blocks?

What's the distribution of clock? So you've made some things easier which is your stage to stage passing of your analog signals but you've made some other things more difficult.

55. Example (2) 12bit Pipeline ADC Floorplan

So, you could do another strategy which is, you do this. I call it an up-down floorplan. I've kept all the amps

together because that means I could put them, share a lot of the wells together, I put all the capacitors and switch together because that makes my clock distribution simpler which is quite challenging.

The uniformity between stages is maintained. The low-res part is a long way away from the high-res part but the scaling of stages becomes more difficult because you end up with the wrong aspect ratio to scale things. And this routing of going up and down, is more problematic, and that could introduce some unwanted parasitic.

56. Example (3) 12bit Pipeline ADC Floorplan

So another strategy is go to the other way of doing the amplifier layout, and the capacitor layout and put them top and bottom and then do it to a left-right floorplan. And this gets in the way of some of the disadvantages, the clocks become easier, because, although they are top and bottom, they're symmetric, you can kind of bring in a clock and distribute it up and down equally to make it mirrored.

But it still has some problems. Often there's never a perfect floorplan. Each of them is a tradeoff of one problem for another problem. And that's why one strategy like I use the U one, in 0.13µm, worked perfectly fine. I used the other one layer in 40 nm, works perfectly fine, but I had two sets of different problems. I decided one was the better strategy for the given technology that I was working on than the other.

57. Handling Density

So, handling density, this is a little bit of a repeat so I'll go through it quite quickly. Basically, if you have a sensitive circuit, you should really fill it by hand, you should observe the matching, you should exclude the auto fill, you should use dummy as shielding for sensitive nets, and you should check the density in DRC with several different starting places so that when you hand it off to somebody else, it's unlikely to trigger at the auto tools or some rework is needed.

58. Chip/Macro Level Considerations

59. Analog \macros/AFEs in SoC

So at the chip and macro level, this is from our 90nm power line generation, we had 2 bands, one would work with 0–30 MHz, one would work with 50-300 MHz. For a necessity of a system on chip, we were given one edge to associate our design with. And the reason, we were given one edge is that we had a certain amount of pins we needed to achieve it and there was no other way we were going to get the pins associated with it.

We chose to separate the bands, we chose to put each of them in their own deep Nwell. We had a common set of references that we put in between. We had to consider what the gradient effects of the thermal, this line drive here, takes 1W when it's on, so of course it's going to have some effect on its neighborhood. And many factors like that have to be considered.

60. Clock/Digital

One of the problems with that block was it was a 5mm wide block and the PLL for the system was on the right-hand side and we had to route the clocks across one band, through one band, drop it off for that band and pass it to the other clock band. We were rerouting the clocks for several millimeters. and we need to make sure we weren't degrading the jitter, the delay, the duty cycle were not too bad. So we actually, when we did the flow upon the AFE, we made the routing channel specifically for the clock and for nothing else, we made it on its own regulated supply with its own regulated shield and we distributed the buffering across it

to balance the kind of delay, jitter, duty cycle changes of it. That determined a lot of the other blocks limitations, that effect that we needed to distribute the clock across. It was not the kind of thing you could have added in at the end but we added certain picosecond requirements for the jitter and we had one PLL and we didn't want to put down multiple PLLs around and even if we put multiple PLLs down, we've still got to distribute the crystal for the reference clocks, so there's various challenges.

The other one associated with digital is in doing that block, we had to deliver it as a block to the people doing the full chip integration. We have to close the timing of that block and we've got routing, we don't want to let our tools route inside our block so we exclude them. The only way to do that is to put known pins of known functions at the boundary of known timing and then to handle that inside your block, yourself will guarantee that it's going to meet timing from then on in. You're giving them effectively a macro that's got known constraints and you have to make sure you have all the constraints for their tools because they need it in a patchy model and you got the bits and pieces and if you don't have that, their tools can't work. So one of the best ways, is to just place known cells, known digital cells at the edge, copy their models in and then work from that point in to your analog front end.

61. Handling Sensitive Analog

Very simple slide on handling analog, basically shielding is your best thing about trying to avoid cross talk and pick up as we intended to do around our references, around analog signals, around clocks, around data paths. We also made some clean supplies for shields and we distributed decoupling and LDO functions in order to try and keep things clean.

62. SI (Signal Integrity)

We have raised something here that is signal integrity, it's a terminology that most digital back end people are very aware of. By here, they say, "timing's clean without signal integrity, but with signal integrity, it's not so clean". So I wonder whether the terminology is correct. You can look for the terminology yourself in the Wikipedia.

Basically, signal integrity is a kind of influence of the neighboring signals on the signal itself and how it effects timing. Could it introduce glitches? These kinds of things.

It's caused by crosstalk of one track to another track. It's also potentially caused by power/ground noise and reflection noise from the higher frequency stuff, but less so, it's more of the first point.

And this is where I point out, just because a signal is static, doesn't mean that it cannot be subject to SI, you might say, "I don't care about this signal, it's just a PowerDown". It happens once, it doesn't matter if its late by a few nanoseconds, I don't care". But you would care, if it actually picked up a signal because that would mean that something toggled, that other block would accidentally be PowerDown. Just because it's static, doesn't mean you shouldn't care about it from a signal integrity point of view. I'll show you some examples about that.

63. Handling Power Routing

Handling power routing, basically the general strategy is to use the thickest possible metal in some form through your power, and if you have many blocks to tie together that are linked, then go to some kind of grid structure, as with any analysis with of a grid structure, versus independent paths, you'll always find that the grid introduces less IR drop issues, but you have to make your sub blocks fit within it, so for example, when we did the analog front end, we published a grid for the analog front end and anyone doing a sub block who wanted to tap into that grid had to make sure that his blocks use the grid at the same pitch at the same

orientation underneath it, because we didn't want that being fixed at the end because, again it's an awkward thing to fix. These are large tracks that don't turn corners very well. It's not so easy to put more vias into them. If you do, you introduce other IR problems.

I put in here, model bond wire and pin inductances and resistances. I should have also put in there, models of actual pad cells, because I would say, the majority of IR drops happen within the first 20 μm of your power pad, because generally they'll have some structures within them that are generalized for certain current densities and most analog circuitries often work at a higher current density that you are expecting, so you might need to electro-migration enhance those pads, put in more top level metal that comes across because many of those power pads were meant to work with any metal stack. So the first thing they do is to go down, and then they go across in the lower layers, when actually you just want to go straight out in the top layer as wide as you possibly can otherwise you're going to drop some IRs very quickly.

64. On Same Isotherm

Regarding thermal, there's a very old reference there. Basically there's a thermal gradient from any kind of major power source on your chip. The package plays a part in evening this out, the thermal study is not just about the silicon, it's also about the whole environment. But generally, what you want to consider, for many power devices, there's going to be isotherms of same temperature. So if you could put things that are supposed to match on the edge, on the same isotherm, they'll do better than if they're not.

65. Example: DAC + Line Driver Symmetry on Isotherms

So in a real world example of this is a line driver, typically you'll have a DAC to drive your line driver, you'll probably want that to be quite close so we laid out the DAC and the line driver together, we made sure it was down the central point. We also made sure that the DAC was a segmented DAC with an MSB section and an LSB section. The MSB section of the DAC goes down the centre line of the line driver. The way we did our common-centroid is we didn't do common-centroid, we did kind of interdigitating around the central line to get the matching because it wasn't just about thermal, and it was also about IR drop because the power pads came in symmetrical around the centre line here. So many things went into the decisions about what to do for that.

66. Optimising Your Design Flow

67. Layout Driven Schematic

So, I talked a bit about the nanometer design involves a lot more dependence on your choice of number of transistors, how they are included, how they are next to each other, so I'd say, I'd like to quote a term which is Layout Driven Schematic. A lot of the tools are driven to schematic driven layout. If you want to close your design cycle better, then you need to early on in your design process, define the transistors for what they would be like in the layout. So make sure you've got these important parameters for stress, important parameters for well effects, make sure you've thought about what the correct unit size should be, make sure you've included 1st order routing parasitics, use a flow that tries to match these things in A/B/C as part of your layout so that your schematic should almost include them in your pre-layout. So before you go to layout, you should have these in your schematic. And that should be part of your instructions to how it does your layout.

So make your schematic, one thing that will help you on this is to make your schematic like your layout floorplan within reason because schematics have to be read by people to understand the circuit. But it can help to segment your schematic and place blocks and show routing so it looks like a layout. But the real

reason to do this is it can help your LVS and your extraction and a hierarchical level to isolate what are the problems, what are the issues, to be included.

And as part of that extraction is, you try and make sure all your routing is part of your sub block layout. Now it's impossible to do something perfect, just do the best effort you can, because the worst thing is something that is unitized to a too small level. And schematics are meant to be read by designers for the function they give you but as part of handing it off to the layout, it's really useful if you can abstract what you need to so it looks a bit like what the layout is going to be.

68. Layout Driven Schematic Flow

So basically, all I'm saying is expand a bit this block and include in it what you're going to do around the wells, what are you going to do around all the other matching effects. And try to match the hierarchy of your layout with schematic and spend more time in this area and hopefully less time in this area, in order to not go around the loop as often.And, that's what we instigated in our team for doing a 40 mn as an approach,

69. Connectivity & Constraint Driven layout

But now some of the design tools are started to bring through these kinds of features. They're doing it in a more regulated manner, so when you do your schematic entry, you input certain parameters and do this what you call Constraint Aware Layout.

70. Fast Validation of Design Intent

I'm not going to do too much advertising for Cadence, because I think they probably have enough.

71. Set-up the Flow Correctly

So, probably more important than the CAD tools itself, because the problem with CAD tools is they're only as good as the user and they're only as good as the set-up that you're using. So the accuracy of the results of your CAD flow for me is determined a lot, of course by the tools and their capabilities, but more so, by the PDK. The PDK needs to be right, you're totally reliant on that to parameterize everything that is in your design tools. Even with the static PDK, there are many parameters to be chosen around the run time of several of the tools, and we'll talk a bit about that later.

Another one, of course, is user procedures, because the tools can do things a certain way. They may only do it if you do them in a certain order with a certain degree of security, so you need to make sure you are using them to the right degree.

Then I'd say, what in the end affects the accuracy of your results, is the number of users. Because in the end, the amount of debug that has happened on your whole design flow. And if you're the only person to use the technology with that set of tools, that set of PDK, good luck to you to some degree. There will always be someone that has to do that job of being the first person around something, but it should raise a big red flag which says, "be super diligent". And so we didn't know what 40nm was going to cause us so as we did in all our technologies, whether it's a new technology or not, we put down dummy test structures in layout, we extracted them, we hand checked them to make sure we are extracting things correctly. We put down some test structures within our chip to give us some extra information, that hopefully in future chips that maybe could exploit about how design rules were arrived at because with this consolidation of foundries out there, you'd hope there would be more and more characterization information but I have a feeling there's probably less characterization information because there's more diversity of circuits. So no one wants to publish a set

of rules and say "that is what you need to do" because it is circuit pacific. So if you want to push the state of the art, develop your own test structures to understand but again, I wouldn't sanction anyone in breaking a design rule because you think you can. You have to be able to prove that you can.

72. Understanding the Optimization

Understand your optimization around the design tools, so for example like extraction, we totally rely on extraction for many aspects of why the layout performs. Most tools, extraction tools have some kind of reduction algorithms, removing minimum resistance and capacitances, merging parallel resistors, merging series resistors, looking at frequency dependent RC reductions, you can say, "I don't care about things beyond this frequency" and it goes around removing things. Well what they're basically trying to do, is to reduce the number of nodes and components within the extracted netlist to something you can simulate.

But, you also need to understand in what order these occur, because if the minimum resistance removal happen before the merging, because if you set at one parameter, maybe you've got thousands and thousands of really tiny resistors that when you added them together, might make a real resistance you care about, but you set a certain parameter and it's disappeared. Once the extraction tools run, it's gone, then you merge them, it's still gone.

So you've got to understand how these tools are doing things, it's the same thing as you got to understand the layout is just an abstraction of the physical and the physical is having an effect on the transistor.

73. Example, CAD Tool Set-Up Challenges – IR Drop-Extraction

So another example, is be aware that labels are extremely important. A label to the tool is a perfect mythical connection that is coming in. So first of all, this can hide an open, you can have a label saying VDD here, another label saying VDD there, and you can have 2 bits of metal that can not be connected. Hopefully your design flows, make sure you check everything from the outside world in, so you avoid this kind of problem.

The other problem is even with just the extraction like resistors, that's a perfect low resistance path between those two nodes. So when it then goes and does its minimum IR reduction, you happen to have some labels left behind here and there, even if nodes are connected. It's going to reduce an erroneous extraction. These things happen.

74. Some Examples of Layout Issues

So, I'm going to go into a kind of examples. These examples are not meant to shame anyone, they're just meant to show some types of issues, they're not all the issues you could possibly see there. I thought I would choose ones that I found interesting, they weren't necessarily that obvious. Just to make you think, I won't have time to go through these examples but you say, "I could have made that mistake, I better be careful next time". Hopefully some you say "I never make that mistake" because maybe you've been through it before, maybe you're much more diligent. It's meant to highlight that, it's not meant to highlight, these are the mistakes that happen, there are hundreds of more mistakes that happen.

75. (1) Metal Induced Error in DAC Current Source Array

So, first of all, this is the metal induced DAC error. We had an array of current sources for an MSB array of a DAC. They were laid out absolutely identical in metal 1, 2, 3, 4, and 6. But the metal 5 was different crossing every second transistor that was the current source of the DAC. When we plotted, we had a mode where we could bring out, we had a thermometer coded DAC, and bring out each DAC current source on its

own and look at its individual matching. We found a distribution which funnily enough, matched exactly the metallization effect, the effect was a 1% delta. It didn't matter in the end because we had some interdigitation of our individual current sources and in the end that wasn't an issue but it made me wonder, how often this happens without people knowing about it.

76. (2) S/H Leakage Capacitance Issue

This is another example. I actually can't remember the parameters involved so I made some up afterwards, but it's meant to highlight an effect.

So we had a designer who wanted to have a very accurate sample and hold, with 16 bit accuracy. They had a switch that was used in the sample and hold and this signal here had higher frequencies on board, but they wanted it to work near Nyquist, beyond Nyquist for an under sampling application. And the problem is, the capacitance across this switch means that the signal is always present on the sample and hold by the ratio of the capacitance across the switch and the storage cap. And the problem was, that around the minimum size transistor, there was a metal layer and we went up in the metal, before going out in the metal, and there was a whole series of small parasitic capacitances and the extraction tool meant that you had a resistor between each of these and each of these small capacitances. So when it extracted it, it lumped all these capacitors to ground, that's fine. To ground, they don't make a difference to the circuit. Unfortunately, they're not to ground, they're to each other.

Now we caught that because we actually did a bit of math before going to layout where we said "to get this accuracy, what parasitic capacitance can I handle?" And then the first time we're doing the layout, we hand calculated the capacitance and we extracted it and extraction tools gave it the wrong result. Luckily, we hand-calculated it and then we tuned the extraction tools and it gave us the right result. We change some of the parameters around it. I can't remember the exact values here, but that was the essence of the problem.

So just to make you aware of it, you should know from a circuit topology point of view, what are your sensitive nodes? Where can you handle resistance? Where can you handle capacitances? If you don't know how much you can handle, don't leave it to the tools to find it for you.

77. (3) A Precision SC Integrator with Variable Gain (Cap-DAC)

This came from some people in Delft, they were making a precision integrator of variable gain. I've drawn on here, basically they decided, they would go down a common centroid design for all of these to get the best all the caps, to get the best balancing and the best gain step accuracy.

What they didn't realize, is that they were increasing all these capacitances. There were parasitic and unwanted capacitances. So really, this is to summarize, a kind of example of common centroid gone mad

78. Is Common Centroid Always a Good Idea

and they cut up each capacitor into four pieces so they could interdigitate them all, they put them all in the centre here, and they ended having 70-80%, 70% of the total area was in the routing. And because the rooting, all the routing needed to have shields, so then you had other parasitic that came about from it. And the real question was, did the inherent process require you to go to common centroid to achieve it? And in the end, they actually decided that it was too big, it also gave them some errors because they got some parasitics around the loop that came back that actually made the circuit not work.

79. Improved Layout

So they ended up doing something much smaller. It's the same area of actual used capacitors but it wasn't done as common centroid between all of them, it was done as common centroid between a selection of them. And it's just to show you, you can go a little bit crazy common-centroiding everything.

80. (4) DAC SI Issue When Placed in AFE

So, here's an example of the SI issue. We had a DAC, within the DAC there were some signals, there was a signal called CLK, and there's a signal called D[0], and there's a signal called CLKINVSEL went to an XOR gate. When it was a zero, the clock would change, and it was there as a safety net, if there was a timing issue, if you're running this DAC really fast, if there is a timing issue, you want it to be able to just bring ourselves a phase change in the data to clock.

It seemed like a nice idea, just cost us a bit, it'll make the design safer. However, we extracted the DAC and there was no SI issue when we extracted the DAC. However, we placed the DAC inside the analog front end, we put a buffer on the edge so we could have a timing model to go to the digital. But there was 800µm routing which is equivalent to 4K resistors now when it's placed in here.

And basically, we had an SI issue when we simulated the whole analog front end, we didn't have an SI issue when we had it as a DAC.

81. DAC Control Bit Example

Now what did the SI issue look like? Basically, when the clock fired the actual voltage on the signal, it would go up by a few hundred millivolts. And when I say the DAC didn't have a problem, it didn't have a problem because we didn't have that series resistors but we still had that capacitance there that caused that. And then if the data fired at the same time, this would actually cause an error.

Luckily we decided that there's enough security in the timing that data and clock would never fire together and we concess that because we found this out when the tape was actually at the foundry. And we had one of those awful dilemmas which is, do you pull the plug, they've got our money, they've got our database but they haven't started the masks. Who's going to sign off so this can go, otherwise we have to bring it all back and there's going to be a month's delay and there's going to be a problem. There's a problem if it doesn't work, which is even bigger.

So, I can tell you, we didn't get much sleep around that time. But we were able to show with lots and lots of simulation, that we were safe and we fixed it. That was just for the MPW but we fixed it by putting in a buffer inside the DAC. So even the sub blocks should probably have buffers on their edges in order that you don't have extract problems when you're doing your signals integrity when you're going up a level.

82. Summary

83. Advice

So, a kind of summary, my advice to everyone is firstly, I do this myself which is I consider myself a novice in a new technology, even if it's the same geometry, if it's just new, it's a different foundry or if it's a different CAD tools. I should be scared. So be paranoid, predict your parasitics and your transistor partitions in the topology and don't believe your results you get unless it's exactly what you've predicted. If it's better than what you've predicted, that's a wrong prediction. The worst one is obvious, okay.

Decide what's important in your circuit, don't be afraid to rip it up and start again. One of my designers, I think we went through an output stage, I think it was seven layouts of it because every time, the first few

times were because we found it wasn't optimum. The next few times were because we suddenly realized that there was something slightly better. It wasn't because what we had wouldn't work, it was a sudden realization that layout is a spatial problem. It's actually very difficult to predict what the best layout is until you see it. You start to get a feeling for the interactions.

But a lot of people go down the route of, I'll lay it out, it's done, it's kind of working, don't change it. But I'd say, as a rule of thumb, the first layout is rubbish. The second layout is normally pretty good, and probably that is where you should stop, and after the third layout, you're probably just tweaking it too much, right? And you should save that for next iteration for the whole thing, because that's too much time being spent on the whole thing.

You floorplan your power, clock, signal flows early in your designs, modular schematics helps in some of these things. And symmetry is only kind of perfect in repeated block and don't be a slave to the automated tool.

I guess I'm a little bit of a luddite personally, in that I still go to a piece of paper first in my floorplan. I know that there are some good tools coming through on floor planning but again, I quite like this ability to get a feeling for the interaction. I like to draw my own signal flows and be able to rub them out, because most times they're kind of wrong. You need to find this feeling as to what effects your design and how to do it.

And the danger to some of the tools is that, I mean they're very good, but they force you into a flow where you don't get to think because you press a button to do certain aspects. So just be aware of it.

84. Disclaimer 2

Final disclaimers, there's going to be a whole load of other effects I haven't mentioned here, I'm sure lots of you will come and talk to me and say, "you didn't mention that", or "you over emphasized this", or "that's far worse than you said". It will always be an infinite list going forward. As I said, every process, every circuit, everything is different. So it's still your problem. I haven't solved anything. Good luck.

And the final thing is, just in case any of you didn't know what a space invader was, you can now see why a lot of people say "a good layout should look like a space invader".

85. References

And I've left some references in the documentation, they're mainly about matching.

There's a good conference which is the Microelectronics Test Structure Conference which is normally around March or April, if you are more in the foundry of the design side of it, or interaction of the foundry side of it. There's some interesting stuff presented, but there's not so much information about layout out there. And what there is, most of the matching studies are what I call area studies, rather than impact studies of secondary effects.

So hopefully that was useful. Thank you.