ISSCC 2011 Tutorial Transcription
Noise Analysis in Switched-Capacitor Circuits
Instructor: Boris Murmann

---

1. Noise Analysis in Switched-Capacitor Circuits

Okay thank you Velneve for the generous introduction maybe a little preface about this tutorial. So I came to ISCC last year and there was an evening session that was called "Stump the Panel" and there were one or two questions related to switch capacitor noise and switch capacitor circuits. I was puzzled about how poorly answered these questions were. The density of misconceptions within 2 to 3 minutes was just so enormous that I thought this is a good opportunity to put together a tutorial for next year. This is why I am here. Hopefully next time when we stump the panel we will be able to answer the questions a little more quickly and precisely.

2. Abstract

So okay here are my slides. So noise analysis in switch capacitor circuits.

3. Outline

I want to give you first a quick outline of what I am going to talk about in the next 90 minutes. Of course I want to first motivate the subject and once I have done this it is necessary to get us all on the same page as far as prerequisite knowledge is concerned. So I will talk about some introductory material that I think most of you know but I just want to set up the terminology and so forth and in case someone doesn't know you can pick it up there and then the rest of the tutorial is mostly driven by examples to make it a little more colorful and these examples are ordered in increasing complexity. I would motivate this a little bit down the road we start from very simple circuit step up a notch and in the end we talk about complete architecture and how to deal with thermal noise analysis in that architecture.

4. Example: Delta-Sigma Modulator

So why are we here? Well switch capacitor circuits are important. For example discrete time delta-sigma modulators and built with switch cup circuits and there if you design for large signal-to-noise ratios you know this. You need to know how to predict the thermal noise because this is ultimately what limits the performance in such a converter. Reminder to turn off your cell phones.

5. Example: Pipeline ADC

Delta-sigma ADCs are not the only ones, the only architectures where thermal noise is essential in terms of understanding. If you design pipelines you know there is a multiplying deck inside that circuit that is also implemented with the switch capacitor circuit.

6. Example: Pipeline ADC

So here also even if you talk about 8 10 bits you need to be on top of the thermal noise to really optimize your architecture.

7. Example: Sensor Interface

Last but not least sensor interface. This is an accelerometer Example. A delta-sigma loop wrapped around the mechanical structure with extremely large dynamic range and here again if you don't understand thermal noise you will not be able to achieve very high resolution in such interfaces. So when you look through all of these applications of switch cup circuits and you could call these fabrics you will identify some common molecules okay.

8. Charge Redistribution Stages

And these common molecules are really your integrators and game stagers. So in essence no matter what circuit you look at: pipelines, delta-sigma, sensor interfaces, things will always come back to building either good integrators or good game stagers. So in some sense once we understand these molecules so to say we will be able to step up and plug them into the architectures to develop understanding at that level. Going even further down if you look at these integrators and gain stages you find an atom that is part of both of these molecules okay.

9. Elementary T/H Circuit

And that is what I call the elementary track and hold circuit. What is that? Well it is a switch and a capacitor so this is the circuit that we will begin analyzing and probably spending a little more time then you would like on this but this is in some sense good for you like broccoli. Okay this is not the most fun circuit but once you understand that one thoroughly you'll be all set and the rest is pretty much a breeze. So we will basically do this in the reverse order that I just presented we will start with the atom move onto the molecule and then do an example of the fabric.

10. Outline

So let's now go into some of the pre-requisite information that you need to understand whatever pertains to these examples that I will cover.

11. Types of Noise

So the first important point to make when I say noise, noise is a very overloaded term in the English language, everything you don't like you somehow call that noise. So I have to be a little more specific about what I mean. So there is definitely substrate noise. There is supply noise and I call that man-made noise. These are all the things that if I don't make a mistake or if I know how to mitigate them these are actually not fundamental problems. But this is not what my tutorial is about. The next level of noise which is a little more fundamental is electronic noise and even there you fork into two distinct kinds one is related to imperfections in technology flicker noise or one over F noise and then there is fundamental thermal noise and when you look at that the flicker noise even though it is not man-made there are good techniques for getting rid of that for example correlated double sampling chopping and so forth and there was a fully dedicated tutorial on this I think a couple of years ago that just dealt with ways of removing flicker noise. So I am also not talking about flicker noise. So the focus of this short course is just on thermal noise okay. Because the thermal noise ultimately if you use all the tricks in the book this is what's left on the table this is what you have to deal with and this will ultimately determine the fidelity of the electronics.

12. Significance of Thermal Noise

So why is this significant thermal noise? Some of you may think well this is kind of a small thing there. In reality thermal noise is really what limits these high-end systems for example you think about audio. You are trying to get 100dB type fidelity wireless transceivers extremely small input signals that you are trying to amplify sensor interfaces resolving sub-electron like charges for example in this accelerometers and so forth and this is really where thermal noise is a major barrier. And in addition to this since most electronic circuits at least now a days it is about low power design and low power design really means understanding noise because if you don't if you manage noise poorly getting low power is very difficult so in some sense the 2 topics are completely coupled you say you are designing a low power high fidelity circuit it is necessary to understand thermal noise. Lastly the significance of electronic noise has been increasing somewhat with technology scaling and this can be seen by this trivial equation down here. Signal to noise ratio means signal power divided by noise power. Signal power divided in the circuits that we build is proportional to a fraction of the supply voltage squared and as the supply voltage has been going down that means to get the same SNR we have to proportionally reduce the noise in the circuit and hey reducing the noise designed for lower and lower noise requires better and better understanding of the noise. And this is why is we are here and this is probably obvious to you and this is why you decided to spend your afternoon with me but in any case it is good to repeat it.

13. Physical Resistor

Okay so let's get started with some very simple things again this is more of a review. You take a resistor and you apply a volt across let's say this is one k without thermal noise you would get a constant current of one milliamp but in reality this current will fluctuate as shown here. And this is what Johnson discovered in 1928 and of course this was immediately shown that this is fundamentally related to thermal dynamics everything that is in thermal equilibrium with this universe has to have some fluctuations and in the particular case for this resistor what we would like to do is model these fluctuations with equivalent circuit components or circuit guys right. So one way to think about it is it's like having an ideal resistor and a current source in parallel would model these fluctuations okay? So the question is now what is the model for this current that goes into the circuit through this current source?

14. Statistical Model

Well because thermal noise is unpredictable in other words if you take a sample of the thermal noise and you look as fast as you can at the next sample those 2 will not be correlated and that is why what we see in terms of thermal noise in the spectral domain is actually just a wide floor and in particular this approximation is very good because you can show and this has to with lattice constant and so forth in matter that this power spectral density of this white noise, so this is how we model this noise source, drops only about 10% at 2 THz and these are not frequencies that are obtainable at least for today. So for all frequencies that we care about we can say that the thermal noise of a resistor is just wide and it is equally spread across the frequency.

15. Thermal Noise Power

So to be more quantitative what is the value of this power spectral density? Well this is what Nyquist figured out early on it ends up to be intimately related with a Boltzmann constant so 4 times kT is the power spectral density of the noise and that resistance and if you now consider a finite interval of band-width so again this is just part of a reminder to you. You must have seen this many times before. So for example if you cared about signals and the frequency band from f_1 to f_2 and you wonder how much noise power is contained in this frequency band you will have to integrate the power spectral density across this band from f_1 to f_2 the power spectral density is 4kT so essentially the power contained in this band is just 4kT times the difference in frequency. This is a very simple integral but we will see more complicated ones later on so this is a good start.

16. Equivalent Noise Generators

So this is the power and then of course if we want to translate this into voltages and currents we can work with a voltage source or a current source to represent this power depending on weather you prefer a Norton or a Thevenin evident type equivalent okay. So this is relatively straight forward so this is resistors. How about transistors?

17. MOSFET Thermal Noise

So it turns out that is you build a resistor with a transistor in other words if you put a MOSFET into the try-out region it will have exactly the same noise as the resistor so there is no cheating there is no such thing as saying "Hi, I trick Boltzmann and I built this resistor with this fancy transmitter and all the noise is lower." Unfortunately this is not the case so the noise is identical. So if you have a MOSFET with a certain unresistance in the try-out region it will have that same power spectral density 4kT. How about the saturation region? Well here things are a little bit different and it can be shown and this is a little bit beyond what I am trying to do here of course that the power spectral density looks somewhat similar so here I am writing it as the equivalent current in the drain. It is 4kT times a constant gamma times g_m times Δf. So the format is almost the same instead of one over R and you have a conductance here and this conductance is the trans-conductance of the transistor and then there is a fudge factor and that's called γ and you have an ideal MOSFET this gamma would be 2/3 and that can be shown with a few simple calculations. However it turns out in modern technologies and in particular this value for γ is no longer 2/3 and in fact there have been some very curious papers in the 1990's that said, "Hey this gamma will go to hell as long as we go down below a micron γ will go to ten-fifteen and you can forget about analog design and short channels." Fortunately these people were wrong.

18. γ Parameter for Short Channels

And so the γ nowadays so this has been nicely clarified by a paper back at the time Sholten at Phillips that showed that you know even the shortest transistors have a γ somewhere around one, not 2/3 but probably in the 90 nanometer process you will looking at one so if you go down to 28 maybe this goes to 1.5 or something like this but it is pretty gentle of an increase it is not completely detrimental. So this is typically the first thing I do when I get a PDK from a new foundry I put a transistor into the saturation region I look at what the g_m is I simulate the power spectral density and I extract the γ. Why do I do that? I know roughly what to expect and so to do this reality check it is good to run this simulation in particular because the more sophisticated a model becomes the easier it is to fool the Boltzmann constant I've found. So in some of these fancy new models it is possible to get a γ of 0.1 or something someone dials in the wrong parameters into the model. So it's become pretty important to check this and so this is kind of important to remember. These γ's are around 1, what exactly, well it depends, but not .1 or 10 or something like that.

19. MOSFET Model in Saturation

Alright so the model for the MOSFET then is simply the usual model that you know with all of it's capacitance, trans-conductance it has a finite output resistance and then it has this equivalent drain current source that models the noise from the transistor. The important thing to point out is that this resistance R_{out} is noiseless because it just models finite dI_D/dV_{DS}. It is just a modeling resistor this not a physical resistor that is noisy this is just a modeling crux. So there is no real noise associated with this. The noise really comes from this generator again that is 4kT γ times g_m in terms of thermal noise and if you wanted to include flicker noise you would have extra terms here but again I want to point out that I am not going to talk about flicker noise again today because there are techniques to remove this almost entirely and then you are stuck with just this part this is why this is important.

20. Noise in Circuits

Alright so we now know resistors, transistor noise. We are going to stick a bunch of these together and build a circuit and this is sometimes what people get confused. So now you have tens of noise sources and how do you figure out what this means overall? So in order to quantify the net effects then of all these noise sources you have to somehow refer these noises to a single point and combine them and what is a good single point where you combine noise sources? Well it is usually the input or the output right? Often times you would pick the input because you know what kind of detection sensitivity you need at the input of that circuit and that's where you would refer all the sources to or you do it to the output and then somehow knowing the transfer function along the way in your calculation anyway you will be able to refer it to the input.

21. Output and Input Referred Noise

So just to show this. So we will have a bunch of noise sources and we can compute the transfer function for each one of them through the output and assuming that they are independent I can just add the powers in the end to form an overall total noise if I want to refer to the input I can first, for example this is one way of doing it, refer all the noises to the output because this is often computational easy and then divide by the overall transfer functions squared to get the noises back to the input, okay. So when I work with input referred noise so essentially what I am saying is I am assuming that my circuit is noiseless and all the noise is abstracted into this input source, alright? So I am bringing this up because this is something that we will be using later on to push noises back and forth into the circuit.

22. Outline

Okay so we are now equipped with all the little tools and background ideas that we need to understand the circuit examples and as now promised I am now starting first with the atom which is the elementary track and hold which that common to essentially all switch cup circuits. You cannot build a switch cup circuit with this as a sub-circuit and we will spend a good thirty minutes or so figuring out how exactly the noise behaves in this circuit and then as promised we will zoom out and apply that knowledge to bigger circuits.

23. Elementary T/H Model

Okay so let's first of all let's establish some fundamentals about some circuit elementary track and hold. The idea is you take the switch and so this switch is a MOSFET and of course I am abstracting away a lot of things that you may actually have in your implementation such as clock wood strapping and God knows how you drive this gate and this may be a transmission gate but in the end it doesn't matter because all of these things from a noise perspective they just make a certain resistance okay when they are on and when they are off I will approximate the resistance as infinite even though that's not true it is actually a pretty good assumption I will comment on that later on. So my abstract model for a circuit of this kind no matter how you build this switch it's just a resistance that is isolated from the capacitor through an ideal switch that is clocked. Think about the behavior of the circuit you apply continuous time input for example when the clock is high the output will try to track this input and when the clock goes low you will freeze the last value on this capacitor so this is how this circuit roughly operates and the time constant RC of the switch and the capacitor of course has to be carefully chosen.

24. Step Response

For instance when you apply a step for example a step that is synchronous with a clock so it could be that a switch cup circuit input is just driven by another switch cup stage and so you apply a pulse at the input and so the output will settle with a finite time constant that is represented by the RC of the circuit and so one thing always to know as a designer is "How many time constants are you settling?" This is an important piece of insight and that will depend on how much position you are demanding so at this point so let's say I am starting to settle here arbitrarily defined as t equals zero and then I wait half a clock period I will have a residual error here depending on the precision that I would like to have in my system I need to wait longer or shorter and that is captured by the number of time constants that I need to squeeze into this time interval. So these are numbers that every designer should know if I want settle to one percent precision I need to wait about 4 or 5 time constant if I'm asking for point or one percent I need to wait for ten times constant before this precision settles in. And I mention this particularly because this number N will pop up throughout this tutorial because a lot of the noise behavior of the circuit will depend on the choice of this N so for the time being you should remember this is going to be between 5 and 10 in any reasonable circuit that has some precision requirements.

<u>25. T/H Noise Analysis</u>

Okay so now let's talk about noise so fortunately this circuit is simple it has one noise source the noise source. The noise source is associated with this resistor and as we just discussed it's voltage spectral density, voltage squared spectral density is 4kTR at Δf so I can pencil this in as a noise generator here, the resistor in series with it and then I have a switch that switches this circuit to a capacitor or not depending on what clock phase I am in. So what I would expect to see at the output is the wave form that looks something like this when the switch is closed the noise source will try to wiggle around on the capacitor. Basically you have a voltage source here that is noisy. Think about this as wiggling around you have an RC circuit that somehow filters this a little but the output should definitely be wiggling cause the filter doesn't completely remove the noise just some part of the noise goes through the filter so I should see some wiggle when the clock is high when the clock is low at least in my abstraction the switch is open and the voltage on the capacitor cannot change so I just freeze whatever value I had here and I hold it until I turn this noise source on again. So the question is what is now the noise of this output and how should we talk about this so if you just say what is the noise at the output that is an incomplete question because you would actually have to tell me what do you mean by noise? How do you want to evaluate it? In particular are you interested in this signal as a continuous time wave form, in other words are you looking at it constantly or are you actually just taking discrete time samples of it or what do you mean noise.

<u>26. T/H Noise Analysis</u>

So for that really the main thought is that since we are talking about switch capacitor circuits the way we want to look at it is to evaluate the noise in terms of its discrete time samples and the reason is that when we turn the switch on and off this is the time when we track the signal and then we have taken a sample of the signal and this is basically the phase in which we care about the noise and we care about its discrete time value. I actually couldn't care less what is going on with the noise in this interval because this doesn't contribute really to my discrete time sample that carries the signal. So the convention for this tutorial and this is how I am driving it is that I am really just computing the noise in the discrete time domain and taking discrete time samples of this process where the noise is switched on it's frozen and off and on again and so forth and these samples in between is what I am interested in.

<u>27. T/H Noise Analysis</u>

So once I say this is what I care about so think about this someone gives you a bunch of random numbers, this is really what it is, a bunch of discrete random numbers, with this index n running from whatever zero to a thousand. The properties that you now interested in as a designer are twofold: one is you want to know what the variance is so basically shown here as a little Gaussian. So every sample you take will land somewhere else on the Gaussian because it is a random process some will land at the mean some will land off and so forth so you can assign a variance to this process of these samples that you are collecting so that is one piece of information and that is very important. The second piece of information can be very important is the spectrum of these samples. In other words you want to know how the noise is distributed in frequencies just like your signal in the discrete time domain has a certain spectral shape. There is more energy at low frequency or high frequency and so forth. You ought to ask how is the noise distributed in frequency as well. Is most of it at low frequency, is most of it at high frequency? What is going on, right? So we need to know the total power and we need to know the spectrum and once we have that we know everything about this noise sequence. So this is what I will work out now and I will begin with a variance because that is somewhat easier and then we will talk in detail about the spectrum.

28. T/H Noise Analysis

So the variance is easy to calculate once you remember Parseval's theorem. Parseval basically says the variance of the samples of a noise process has to be equal to the power spectral density of the noise process integrated over all frequencies and so this is all you have to do essentially to find the variances in these noise samples. So variance is really just a statistical term for the RMS squared noise voltage in this case. So what I need to do to find it is according to Parseval. I take the power spectral density and integrate it from zero to infinity. What is the power spectral density? Well it is whatever I have from the resistor $4kTR\ \Delta f$ this is what goes into the circuit and then that's filtered by the RC. What is that? Well it is one over one plus $j2\pi fRC$ and so I can carry out this integral it just takes about 2 lines of math and this is of course a constant and it goes outside the integral and this integrated gives you one over 4RC. The 4 goes away conveniently surprisingly the R also goes away and the C remains and the final result is very famous it is kT on C. So this means the variance of these samples that I am taking is equal to kT on C and what is puzzling at least at first glance is how is this possible? The noise comes from the resistor but the variance doesn't even depend on that that resistance. This is bizarre it depends on the capacitance even though the capacitance is not contributing any noise. The capacitance is clean.

29. Effect of Varying R

Well the reason for this is that at least from a circuit perspective this is how you can think about it when you change the resistance the power spectral density changes. So 4kTR for larger R you get a higher power spectral density but your corner frequency goes down so the integral of this green curve and the integral of this blue curve are exactly equal and R actually drops out in the final result and here I am actually plotting the running integral of the green and the red curve. The only difference between them is that they converge at different frequencies the final result is the same and this is kT on C and it's a good number to remember as a designer always have to have it in the back of the head if you make the Cap 1 pF kT divided by 1 pF and the square root of this is 64 μVrms and this is what this is. So I used 1 pF as an example here. Another important piece of information from these plots to remember is that these noise integrals they tend to converge about ten times beyond this corner and so you see this green corner is around here and the integral hasn't quite converged yet here you have to go pretty much another order of magnitude in frequency to converge. Why is this important? Well it tells you how far out you have to consider noise and the integrations and so forth. Once you do calculations and simulations and so forth and we will reuse this factor of ten a little bit later on.

30. Alternative Derivation

So this is how circuits folks would derive it. If you talked to someone in physics they would derive this in a completely different way they would go and say "Hey I remember the equipartition theorem!" which says that in general, forget capacitors whatever you have anything that can store energy in the universe, it could be a spring mass damper system, it could be whatever you have, an inductor, anything that can store energy and

Is in thermal equilibrium with the universe more or less holds an average noise energy of kT on 2 and so in order to apply this to figure out what the noise variance on our capacitor is you have to figure out what the capacitor energy is equate that with half kT and that will give you an indication of what the voltage noise is.

Or for other systems I don't know the velocity noise, the displacement noise and so forth. So for a capacitor clearly the energy stored is one half CV squared and according to the equipartition theorem that is equal to half kT so you see where this is going it's just for v_{out} square bar kT on C. So no need to integrate anything or something like this. This is really fundamentally rooted in the way the universe is constructed. You don't need to be an electrical engineer to come up with kT on C. Why am I talking about this? First of all it is kind of interesting to see it second of all I will actually use this later to help us simplify an analysis in a real circuit.

31. Does kT/C Noise Matter?

So this kT on C business. Does it matter? Let's just plug in some numbers and this table is generated of course assuming that my circuit is this RC little widget that I have been talking about so far and putting in the one Volt sinusoid and then I am asking this question at the end of the day what is the SNR of the samples that I am taking? So we know the noise is kT on C and I know what the power of a one volt sinusoid is. It's just finding the ratio and I am expressing that in dB and then I calculate what is the required capacitance to meet this SNR. So here you see that basically a very steep trajectory in capacitance level that you go through when you talk about 20 dB versus 140 dB if that even exists so if you say I am going to build a system of 20 dB how should I size the capacitance to meet my noise specifications? The answer is don't worry about it. Because any capacitor that you can build will be bigger then this we are talking about aF here or something like this. This is why when this week you go to a digital session or something like this no one talks about thermal noise. Because these guys compute with SNRs of I don't know what maybe 20 dB they don't care. On the other hand if you build very high fidelity systems you actually start wondering how I am I going to build such large capacitors and this is when you resort to techniques like oversampling and at the end of this talk I will actually show you how you manage to build very high SNR potentially without using such large Caps. And then there is of course middle ground you know these are the pipe line guides and so forth you build a ten twelve bit circuit you end up working with a few pF of capacitance and that tends to be manageable but none the less you need to know exactly how to get to these numbers. So as a disclaimer again in reality your circuit will not exactly behave like this because this is just a noise for a very simple RC circuit we will see how these expressions evolve as we get to more complicated circuits but this gives us a ball park. This gives us the kind of the unit calibration that we are interested in as engineers.

32. Spectrum of Noise Samples

Okay so now let's answer the second question that I posed earlier what is the spectrum of the noise samples? So I know just to sum this up again I know I collect a bucket of these samples and I try to find what the variance and the variance is kT on C but now I want to know how is this noise power distributed in frequency so I have a discrete time process that of course has a spectrum I just take the discrete Fourier transform and I should be able to tell you what is the spectrum of these noise samples? So the intuition for this is as follows. So imagine looking at this sample and determining what it's noise is if the RC time constant of the circuit is very small that means the noise source itself is very successful of wiggling the noise back and forth in the capacitor in the allotted time. So I am waiting half a clock cycle if within this half clock cycle the noise source can wiggle the noise very quickly because my time constant is very small. The correlation between this sample and that will be very small. Why? Because I am doing a lot of random experiments in between I am rolling the dice a hundred times. Nothing will be related to each other. On the other hand if the time constant is very large that means the noise first of all will not move so hectically here it will be a much smoother wave form and that means the history from the previous sample is still contained in the new sample and that means correlation and correlation always means that you have some energy pushed towards lower frequencies and so in other words what I expect if I did a very thorough analysis of this which I will do just for completeness that the spectrum should depend on the time constant of the RC circuit, okay.

33. Transient Noise Simulation

If you don't want to do the math and just play with it a little bit you could run for example a transient noise simulation. We will talk a little bit more about simulation later on but here is an example that just confirms this intuition that I am talking about. So here I am talking about N equals ten means and just to remind you N equals ten means a lot of time constants fit into this half clock cycle so the noise can freely move many times up and down from sample to sample. That means the noises should be uncorrelated from sample to sample so the samples are these plateaus.

34. Transient Noise Simulation

On the other hand if I dial in something like N equals one well you have a very slow and steadily moving noise behavior between the samples and that will naturally force some correlation between the 2 samples. So if you have access to a transient noise simulator you can play with this to build this type of intuition

35. Calculating the Spectrum

But now let's do the math and again this is something you will not use this at all as your daily life business as an engineer. You will hear this once and say "Ah this is interesting" and then you will tuck it away and never touch it again. So this one time will be now, okay? So how do you actually compute the spectrum? So the spectrum comes from the Fourier transform of the auto-correlation function and this is yet another theorem that I am throwing into the mix here this is a Wiener-Khinchin so if you know the auto-correlation of this three time process you take the Fourier transform and you have the spectrum. So the derivation that I am showing now is therefore focusing on finding the autocorrelation of the noise at output and then I am taking the discrete time Fourier transform of this to be able to plot the spectrum okay. So there are 2 steps.

36. Analysis

So the first step is how do I find the auto-correlation at the output of this sampling circuit? Well it is not too complicated it turns out. What is the auto-correlation of the input? This is my resistor noise. It is white and so the autocorrelation function is a delta function. Why is it a delta function? Well white noise is only correlated with itself and that's why it is a delta function at the origin in terms of auto-correlation but now I am sending this auto-correlation function through a filter that has a certain impulse response so this is the impulse response of the RC circuit and how do I get the autocorrelation at the output? Well now I need to solve this convolution and this takes about 3 lines of algebra and what you get in the end is this it's the auto-correlation at the output interestingly is kT on C times e to this τ hat divided by τ. τ hat is really the difference in time between which these samples are taken and just as a reality check this result makes complete sense because for τ hat equals zero that's just looking at one sample at the same time that is actually not the auto-correlation that is just a variance and we already know that is kT on C. So this is consistent with what we know it then basically says this auto-correlation dies out for other values of τ so what we need to inject here in this τ hat is the length of the noise process that we invoke to disturb the circuit and so this is then the covariance of samples separated by n clock cycles at the output. So I can basically this is a general result it not only gives you the covariance of samples in adjacent clock phases but also the one next to it and so forth so let me now

37. Analysis

take the Fourier transform of this. That's simple so this is the expression so you evaluate it again 3 lines of algebra later and you finally get what I wanted and this is the expression for the spectrum and this looks pretty complicated. Luckily it has kT in C and so forth it has a bunch of cosines and things like that. So there is not much that you can do with that other then plotting it and the key parameter in this expression is my N and again N is the number of time constants that fits into half clock cycles and that is when something interesting comes out. So according or confirming our intuition actually for N equals one, again this means the arch time constant, not much time for the noise source to clean up the previous sample or to disturb the previous noise sample. There is a ton more energy at low frequencies or power where as you increase N this becomes wider and wider and wider and for N equals 5 well for an engineer this is the right spectrum isn't it. So in other words if your time constant is small enough such that the circuit actually settles which we required earlier by looking at it we said: "Hi, five to ten that is what we are going to use." This means the output spectrum of these samples will be wide and wide in particular implying that the power spectral density must be 2 over f_S times kT on C. Why? Because this whole area under the box we know is kT on C and the frequency range from here to here is f_S over 2 all the power in a discrete time process is contained from zero to f_S over 2. I just divided by this range which is 2 over f_S and this is my power spectral density of this process in the discrete time domain, okay? So just so that you believe me that this is not just mathematical nonsense you can go home tonight and fire up your simulation tool.

38. PNOISE Simulation (More Later)

If you have PNOISE available and confirm this and low and behold I just did a PNOISE analysis I will explain later how this works and how to set it up and so forth. But if you run PNOISE with different values of N on this RC circuit and you look at the power spectral density that you get out of the simulator you will get exactly the same curves that I derived earlier. This is by the way a very nice hello world simulation example then if it's the first time you've run a PNOISE analysis try to duplicate these curves if you can't you are not using the right set up for example because this is fundamentally how it should behave. But again the bottom line is we don't care much about this coloring in practical circuits because we will always work with large N that will make these noises look white. Okay.

39. Noise Aliasing Interpretation

So one thing that has been bothering designers quite a bit is that power spectral density in the discrete time process it seems quite high so this value of 2 over f_S divided by kT on C is high because when you compare it to the noise PSD of the resistor remember this is where you started this is the guy who makes the noise it has this power spectral density. If you look at the ratio of these 2 so the PSD of the samples and the PSD of the resistor interestingly you can divide this out and it ends up being N. In other words if my circuits settles for N time constants the sampling operation will increase the power spectral density by N so the discrete time process has an N times increase power spectral density compared to the continuous time noise source and you can interpret this in whatever way you want but the most popular way of thinking about it is that this increase is really due to the fact that we are compressing the frequency spectrum between zero and f_S over two and all this noise aliases from higher frequencies into this band.

40. Noise Aliasing Interpretation

So here I stole a diagram from that is in one of Ken Kundert's white papers where he nicely draws this. He draws this 2 sided but you can also imagine this being a one sided thermal noise spectrum with just the resistor with the capacitor and the continuous time domain and once we sample you basically slice this into pieces and due to aliasing these pieces land on top of each other and this is actually also the hand wavy explanation as to why the noise in the end is wide. Well you take so many pieces of different shapes and most of them are actually pretty white and if you stack them on top of each other you end up with this box that is basically white noise. We know the areas kT on C so the power spectral density has to be 2 over f_S times kT on C. So just to contrast this with my analysis this is what you find in text books and so forth you combine all this stuff the noise is white don't worry about it. But what I showed you is the mathematical proof and again this is the thing you want to tuck away you don't need it any more just know that it is white.

41. Charge-Redistribution T/H

Okay so now we have geared up to a graduate from discussing the atom to the molecule so I now want to go in and say okay let's apply what we learned to a more sophisticated piece of circuit which is a charge-redistribution track and hold and now things of course are much more complicated but with the knowledge that we obtained we are able to dissect this into small manageable pieces and combine them in the end to end up with the complete expression for how this circuit fares in terms of noise performance. So first of all let us just quickly review how this circuit works. So there are 2 phases as always as there are in most switch cup circuits. In phase one and one is on and 2 is on and so the input voltage appears across the sampling capacitance C_S. At some point in time there is an early clock transition that opens this bottom plate switch at this instant the charge at this point is frozen then this switch is disconnected Φ_2 goes high which turns on this switch and the charge that was acquired here in the first phase is redistributed to the feedback capacitor by virtue of this virtual ground of the amplifier. So why does the charge have to go there? Well you connect this capacitor grounded here the amplifier makes a virtual ground no charge can remain here all of the charge has to go on the feed-back path. So the way to analyze the circuit in terms of the signal transfer function and I do this just to emphasize how that is properly done so that we can do the same thing for the noise is to properly analyze the circuit in terms of its charge behavior. So this is an important piece of information to remember when this switch opens the information on this signal is contained on the charge at node X. So in some sense in the sampling instant you go from voltage to a charge and then in redistribution you take this charge and you turn it into a voltage, okay.

42. Charge Conservation Analysis

So the way to analyze it and you know this looks trivial during Φ_1 I acquire a charge it is just $C_S V_{in}$ during Φ_2 the charge is just on the feed-back capacitor and that defines V_{out} so the charge is $C_f V_{out}$ and the charge at this virtual ground note cannot change. So that is called charge conservation. So these 2 charges have to be the same. I plug in the above results and that means I can solve for the out of the function of the in which is just a ratio of the 2 capacitors. What is important to remember though even when you see this end result even though you don't realize it there is actually a middleman in the game. So you go from voltage to charge and then to voltage. It is actually not a voltage-voltage transfer function it is a voltage to charge, charge to voltage and I emphasize this because when we now look at noise we have to be very careful about counting the noise and this charge that is being re-distributed because that is also where the signal is.

43. Noise Analysis During f1

Okay so as promised we are going to dissect this circuit now into smaller pieces because analyzing it all at once is just a big mess and we will understand nothing. So the first thing I will do is split the analysis between Φ_1 and Φ_2. So let's first look at the noise that is introduced in the sample phase and then we will look at the noise in the re-distribution phase and then we will somehow combine the 2 and that will help us get an overall expression. So thinking about Φ_1. So here is what is really going on. My amplifier is actually reset with a big fat switch at the output which means it is basically doing nothing I may as well delete it because it's output can't move so it does nothing here to enforce a virtual ground and so forth. It is really not doing anything I may as well remove it. I have all of these switches closed. So this one and this one and the sampling instant is really defined by this bottom plate switch opening. So this Φ_1 going low first. This is again when the signal is acquired in form of charge at the node x. So the right question to ask and this is very critical and a lot of people are highly confused about this. The right question to ask is how much noise charge is introduced in the sample instant? This is the right question to ask because the input signal is acquired in form of charge and so I ought to be asking how much noise charge is acquired at the same instant. You cannot just analyze the input voltage that is sampled on C_S in terms of charge. That is completely wrong and it leads to really bad results and I will show you why in a moment. So having said this, how do I find the noise charge at node x? This is actually pretty tedious because if you think about it even disregarding this because it is a very small resistance that maybe doesn't make much noise. I have at least 2 noise sources here and both noise sources have an interesting or a multipolar or zero type transfer function when it comes to referring their noise generators to the respective charge here. So I would have to write basically the mechanical way of doing this would be put a noise source here, calculate the transfer function from this voltage to this charge, integrate from zero to infinity and then do the same thing for this and potentially do the same thing for this and repeat. Turns out that is a page of algebra and once you have done it you will realize, man this result looks so familiar! I should have been able to just eyeball it. And this eyeballing can actually be done with the help of the equipartition theorem.

44. Equipartition to the Rescue

So what is nice about this circuit is it is actually purely passive and purely passive circuits are in thermal equilibrium with the universe. This means equipartition must apply. So in other words if I wonder how much noise charge I am acquiring here, all I have to do is identify what is my energy storage variable here or what is my degree of freedom in this system? And this degree of freedom will have a noise energy of half kT. So the degree of freedom clearly is the total noise the total charge at this point and how do I compute this? So I need to somehow compute in terms of the energies if I knew the charge energy is one half q squared divided by C. But what is the effective capacitance that I have to consider? Well it turns out that it is just the sum of all the capacitances connected to this node. Why? Well because from an energy stand point it doesn't matter whether or not you have a series resistance in this capacitance. If you want to put charges on this capacitance energetically that is the same thing if there is a resistance or not. You have to expend energy to put charges on the capacitance. So this definitely must be true from an energy perspective and we now know that if this is really the energy stored in the form of charge at this node that must be equal to half kT and so now again I can solve this for the noise charge of this node and it ends up being kT times C_S plus C_f plus whatever parasitic capacitance you have hanging here and again this is the point where after doing 3 pages of math you are wondering there has to be a better way of doing this and this is it. So a few words about this, so I will show this later a little more explicitly. This is a term that can be highly detrimental. In other words if you build an amplifier that has a huge parasitic capacitance at this virtual ground note it can really destroy your noise performance. I have a little anecdote about this. Somebody hired me for a consulting job and they said you know we have a big sub-straight noise problem in our chip can you look at it? And a few days later I figured out that they don't have a sub-straight problem they don't know how to compute thermal noise because their parasitics at this node were so large relative to the really tiny sampling capacitance they were using to minimize power and so forth that essentially the thermal noise that they predicted was really wrong. So this is really an important consideration in making a low noise switch capacitor circuit.

45. Noise Analysis During f2

So now let's move on to phase 2 and the first thing to realize is what will happen to this noise charge that I just computed in Φ_1? Well it will be redistributed just like the signal charge there is no difference so I had some signal charge here and I explained earlier well this gets pushed to the feed-back faster. Well guess what happens to the noise charge? The same thing so whatever noise charge we computed if we want to refer to the output we just drop it into C_f and that will make an equivalent noise voltage at the output and I am showing you this a little bit later when we combine all the results. What we really need to do and this is going to be a lot more work is to figure out how much noise is now including or injected from this other components that come into the game so definitely I have some noise charge that will go to the output from Φ_1 but in Φ_2 I have some noise processes running as well now and I need to account for those. So for instance the switch that does the redistribution while it has on-resistance., this means that it comes with a 4kTR noise source here, and I have an OTA here and this OTA will be noisy and so both of these elements will affect how much additional noise is now introduced in this output sample during Φ_2 and so my goal now is to analyze this and what I need for this obviously is I have to make some assumptions about this OTA. So I need to have a model for the OTA. I need to talk about what kind of noise it generates and what I decided to do here is to use the simplest possible model which is a single stage OTA and once you understand this it is pretty easy to go for a more complicated OTA and I will give you a reference in the end essentially the expressions remain about the same with a few additional factors. But this is kind of the basic entry to understand what is going on.

46. Single Stage OTA Model

So my OTA model therefore as a said single stage so therefore I assume

this may be a fully differential circuit it may have a common node feed-back and so forth. But at the end of the day if everything is biased upright I can draw a single ended model of it and what is common to all of these single stage amplifiers then is there is a transistor then that is really the transconductor that you want and that of course has its noise $4kTg_m$. But there is also noise from current sources and things like that. I have to bias this guy somehow. So for example I have a PMOS up here that gives me $4kTg_m$. I didn't want this and actually don't need it but I don't know how to make a current source otherwise. So the way that I can model this is as follows. I can say well these 2 basically just add and what they do is give me an equivalent output current source that I just connect to the output of an otherwise ideal transconductor and the power spectral density of this current is just $4kT\gamma g_m$ of the NMOS and $4kT\gamma g_m$ of the PMOS and to write this a little more elegantly it is usually done like this you pull out the g_{mn} and so it becomes $4kT\gamma g_m$ times α where this α is a term somewhat larger than one. So α equals 2 if these 2 transistors have exactly the same g_m and you will see this all the time if you are a good circuit designer what you will try to do if you can is make as little g_m as possible with this upper transistor and you can do this for example by using the smaller device with larger gate over drive. It costs a little bit of head room but it will help you to minimize this output referred noise current of the OTA. So this is just a aside as a design consideration. In any case this is something we can now plug into our circuit.

47. Noise Analysis During f2

And that is done here. So now we basically at least have the two main sources plugged into this and we can use this now as a template for further analysis. So the first thing I want to assume now this is critical, if you don't do this things get very messy, is to say that this R_{on} more or less it should not at all affect the dynamics in the circuit too much, why? Well if it did I probably did a poor job as a designer. If my circuit is limited in speed and settling and so forth by the resistance of the switch rather than the finite transconductance of my amplifier I am basically beating a horse with a stick here. It is easier to make very small on-resistance then to make very large g_m and so most of the time and you can of course deviate from this were it is appropriate but most of the time and in particular in this analysis it is assumed that this on-resistance is so small that it doesn't affect the dynamics of the circuit very much. So assuming this the circuit becomes a single pull system and I can capture the single pull by basically just identifying at the output of the OTA an equivalent resistance and an equivalent capacitance and so I have a single pull system with that time constant that represents the behavior and the dynamics of the OTA. So let's do this.

48. OTA Time Constant

So how do I find for example the equivalent resistance at this node. Looking into here this is relatively easy. So consider connecting a diode connected transconductor. If you connect this output to the input this is a well known result. You will see an impedance of one over g_m right like train and gate of a transistor connected together. This is the same thing. I don't have these connected together instead I have a capacitive divider in between C_f and C_S and $C_{par.}$ And this divider ratio is usually written as β. So just C_f divided by all the other caps hanging at this node and so with this division my transconductor delivers proportionally less current according to this β and that is why the resistance is not one over g_m but one over β times g_m, okay? This is a very easy way of just eyeballing what this equivalent resistance is. What is the equivalent capacitance, well definitely I see whatever load cap I have hanging here and then looking back here I see some capacitance and what is that? Well since I already know β I just invoke the Miller theorem. So I see C_f times one minus β. Remember Milller Capacitance gain across the capacitance and this time the gain is less than one and one minus β is a number smaller than one so you see only a fraction of C_f looking this way depending of what β is of the feed-back factor. So now I want to know what the time constant of this amplifier it is. It is just the R_{eq} the C_{eq} of what I just calculated.

49. Output Referred Noise

And that is why I know that the transfer function of this circuit, anything that has frequency dependence, will be just first order within an H(s) one plus $sR_{eq}C_{eq}$. So now I am in business and I can now write transfer functions from both of the noise sources to the output incorporating these dynamics over which I will need to integrate later on. So at the output now what do I see? I see this noise source it goes through the amplifier, the closed looped transfer function of the amplifier. So the closed loop, assuming this has infinite gain, the closed loop gain is just C_S divided by C_f and I have to square that because I am dealing with powers. And then in terms of dynamics this is shaped by H(s). This is just a frequency response of the circuit, a single pulse system. It's given by this equation that we just derived. Then I have the second noise source. This current wants to inject something here at low frequencies it is just iR or i^2R^2 because again I have to raise it to the power and then again modified by these dynamics. So now this is a crucial step now because overall this looks like a big mess right. How am I going to develop intuition about this? Well the good news is that one of these terms tends to dominate so what I am trying to do now is compute the ratio of these 2. So on the next slide I am writing A divided by B and I will show you that one of them dominates and we can basically kick out the other.

50. Output Referred Noise

So what is this ratio A divided by B? So it is the resistor noise divided by the noise from the amplifier and you can massage this in particular this i^2 contains g_m. This contains 1 over g_m^2. So what remains from this is 1 over g_m. This R_{on} stays here and then you have a collection of terms that add up to something of about the order of one, so not too different. So you know γ is in the order of one, α is maybe 2 or something like this, $β^2$ is actually C_f divided by C_f plus C_S and so forth. So the product of these 2 is not too far away from one. So at the end of the day the contributions split between A and B as the ratio of the on-resistance and one over g_m again and so this brings me back to my initial argument I would be a pretty bad circuit designer if I don't minimize R_{on} because that is easy I just throw in a fat switch or at least a little bit fatter then my transconductor. By the way I also have much larger voltage then on that switch then on the transconductor. So it is almost always possible to make this A term go away and be entirely limited by just the noise from the OTA you will see this later in the simulation example that I have. You will then have God knows what 5-10% noise from the switch and 90% noise from the OTA but for a hand analysis you will of course knock out that little piece and proceed with a simplified expression that is good for intuition building and simple hand analysis.

51. Output Referred Noise

Alright so now that I have thrown out all this stuff I can actually just go and integrate the total noise introduced by this OTA. Now by the way I cannot not use the equipartition theorem because this is not a circuit in thermal equilibrium once you have active devices in saturation. Those are not thermally in equilibrium. So you have to do this noise integral low and behold it is easier right because we threw away everything. It's just one integral and by the way it is the same integral that we solved through the simple RC circuit it just has a few more constants in it. It is the first order shaped PSD that is basically these terms. So you do this integral and again you see one over four RC coming out of this. These are just constants. When you expand this R_{eq} this has a β and a g_{mn} in it and you cancel all the terms something beautiful comes out and again this is what I mean by somewhat intuitive. Again it is kT on C and it is kT on $C_{eq.}$ It is the equivalent capacitor at the output of this amplifier and you have kT on C times something. So there are some constants involved so it is not surprising that γ shows up, sure α shows up. That is the excess noise that I get from the current source and then one over β shows up. So the total output noise during $\varphi_{2,}$ the noise that is introduced during φ_2 depends only on these parameters γ α which is your amplifier topology technology, the total note capacitance and then your feedback factor. So this is a pretty beautiful result and you will see this again and again and again. The total integrated noise of any circuit like this even if I threw in a 2 stage amplifier or an N stage amplifier or a cast ordered amplifier it always kT on some C time a few constants and this always shows up. Okay.

52. Adding Up f1 and f2 Noises

Now we are in a position to just put everything together. Why is it easy to put this together? Well I am combining now output noises from physically independent noise sources so the noise charge that I sampled in φ_1 came from some switch that hangs on the bottom. That noise that I introduced in φ_2 came from the OTA. Those are physically independent noise sources. And hence I can just add the powers. So this is what is done here. So how do I do this? This was the frozen noise charge at the virtual ground during φ_1 kT times the sum of caps and if I want to refer that to the output voltage is charge divided by capacitance and if I have charge squared then I have to divide by capacitance squared and so this is my output referred noise from φ_1 from the sampling instant. This is the charge I froze there. And this was my OTA noise from φ_2. I just put them together. So my total output referred noise expression is this. If you want to input refer it there is nothing easier then that. You just divide by the closed loop gain squared of the circuit which is C_S divided by $C_{f.}$ So the input referred noise become this and this is again a good expression to look at a little bit and to step back and so the first thing to note and again this is pretty important the noise from the sampling phase input referred is not just kT on C_S. you will find this in almost every text book there is but this relies on the approximation that no other capacitance is connected to the charge conservation node. In reality it is kT on C_S times one plus the ratio of other caps connected to this node relative to C_S. As I mentioned previously if you do a bad job with your layout or designing this circuit somehow skewing it into weird direction such that the parasitics are bigger than the sampling caps this noise can grow out of bounds very quickly. So you have to be careful about this. Also to my surprise this is a term that is also often dropped by simple analysis in papers and text books but it is actually equally large then kT on C_S depending on what your load capacity is and so forth. So basically 2 component sampling noise. Redistribution noise depends on how much cap you have at the output. If you don't load the circuit with a very large cap this can actually dominate and again if your parasitics are large this inner term can dominate and none of these terms are widely available in literature. So this is something to care about when you do your next design.

53. Noise/Power Tradeoff

Alright so now that we see this final result it is time to pause for a second and just remind you and bring you to the initial claims that I had that oh you know noise is really important. This is why.

54. Adding Up f1 and f2 Noises

Suppose I now design the circuit and I compute my noise based on the capacitors I have and it is a certain number that gives me a certain signal to noise ratio. Now my boss comes and says give me 6 more dBs because, I don't know let's leave some margin.

55. Noise/Power Tradeoff

What does that mean? I will have to increase all caps by 4x unless I completely destroy my circuit or the architecture and design something different because the noise is proportionate to kT on C and 6dB means 4x in power. In order to maintain my speed. My boss didn't tell me to drop the speed. So I need to boost up the transconductance to maintain the same speed so I have to make 4 times more g_m in order to make 4 times more g_m in order to make 4 times more g_m I will increase my current by 4 times and I maintain the same g_m on I. Why do I maintain the same g_m on I? Well, otherwise I will destroy my transit frequency so if I go to higher g_m that means that didn't optimize the circuit very well if you go to lower g_m on I, well you pay even a bigger price. So the bottom line is that 6dB or one bit if you think about what that means in terms of quantization error. Quadruples power dissipation in a noise limited circuit and this is something that a lot of people especially at the higher levels do not understand. You cannot over design randomly by 6dB. You cannot because it kills you in terms of power dissipation. So unless you figure out really what you want don't even start, okay so there is no joking around with 6dB back and forth. It is really, really steep out there.

56. Noise in Differential Circuits

Another reminder to step back. All of these calculations I did where with this little abstraction with this single ended circuit. What if you have 2 half circuits? Well, when you have 2 half circuits the noises in the 2 half circuits can be treated independently and so the noise power is essentially doubled, okay? And the signal power actually increases by 4 because I have twice this way and 2 squared is 4. And so you can argue and you see this in text books that differential circuits have a 3dB advantage in dynamic range and this is true but you are also investing twice the power, okay. So you are building another half circuit. So the bottom line is really that there is no magic in differential circuit that gives you a real advantage when your power constraint because you may as well put the same power into the single ended circuit to get the same dynamic range so there is no win in power dissipation by going on a differential.

57. SC Noise Simulation

Okay so now we are at a point where I have analyzed the circuit in great detail and this is where it is always a good idea to throw this into the simulator and see if this is really the reality right. I mean I could have told you whatever and it is not reproducible by any means and so I want to show you now a few methods by which you can actually verify exactly what is going on in this circuit and I used this of course to check my answers here and also to introduce you to the various methods. So the first method and this is a very low cost one because you can do this with any simulator out there that has a dark noise capability is to basically do the same thing that we did in the hand analysis. You put the circuit into φ_1, You run a dark noise analysis. You integrate the noises. You get the variance out of this. You put the circuit into φ_2 state. You integrate all the noises and then you take a piece of paper and you add the 2 numbers and that is your total noise. Okay so that is one way to do it and I will show you that.

58. SC Noise Simulation

A slightly more sophisticated if you are familiar with periodic steady state simulation PSS you find a periodic operating point of your periodically switched circuit which is a switch cup circuit for example. And then around this periodic operating point you can inject these small noises using a PNOISE analysis in the frequency domain and when you do this one little advantage and I guess this takes out a little bit of labor and probably makes it less error-prone is that in this analysis the circuit is automatically clocked between the 2 phases and the noises are superimposed automatically by the simulator so the end result you get in the end it is just a total noise. There is no adding things up on a piece of paper. The third method is transient noise and that is somewhat intuitive and in some ways the most physical way of simulating the noise because really what you do is the simulator injects a thermal noise generator in each transistor and you just look at the transient wave forms. It is also good for debugging you can kind of see where the noises go and what they do and so forth. So I will show you now a run through the circuit that we analyzed using these 3 methods and then I will compare the results.

59. Example T/H Circuit

So just to prove that I really did this you know I actually entered the circuit here, here is my schematic. No special numbers here so 100 MHz sampling frequency. My α is 2 which means the noise from the current source in my signal stage OTA equal to the noise from the transconductor divides itself. The γ is one in this technology that I used C_S C_f C_L I dialed in certain numbers and I didn't bother putting much parasitic here, if you could do that it doesn't really make a difference. So how is this circuit designed first of all? This is actually pretty important to obtain reasonable results so the transconductor is designed such that ten of its time constants fit within half a clock cycle. This sort of what I want this OTA is determining the settling speed here and I want ten of those time constants to fit within half clock cycle the switch is I make 5 times faster and why do I do that? For noise reasons okay and also for settling reasons because I do not want my switches to limit my speed I want my OTA to limit my speed.

60. Noise Simulation(f1)

So then I do my dot noise analysis and the first one is just to find the integrated noise charge here. You may at first wonder how do you do this? It looks a little bit like magic. But it's actually easy. How do you find the charge at a node? Well you tell SPICE to compute CV and CV and you add it up. This is what is done here using a behavioral source CV plus CV and then I divide by C_f to refer directly to the output virtually in some sense.

61. Noise Simulation(f1)

So I do this. I simulated it. Here is the plot I get. This is essentially already this charge referred to the output and now I integrate that and this is what I get 400µV, 406 µVrms. So this is the first number I put on my piece of paper that is the noise from phase 1.

62. Noise Simulation(f2)

I do the same thing for φ_2. I just turn on the other set of switched run my dark noise analysis. I integrate at the output to 266 and now I combine the 2. I combine them not in rms of course. I have to combine the powers and so this is my total output preferred noise. I can now of course divide by the closed looped again squared to get it to the input so this is my end result. So this is how I would do this as a designer.

63. PSS Simulation Setup

So for PSS and PNOISE you need to do a little more work in terms of setting up the analysis and you need to be a little bit careful. So there are various parameters in PSS for example to find a periodic operating point. You need to say starting from which time is the circuit stable. So if you have some initialization in clock loop strapping you have to go through a few circuits before the circuit is really in a steady state and this is where you would specify this. Then there is an important parameter it is called "maxacfreq" that tells you how far out do I need to simulate the noise to account for this folding properly that takes place. So you have to remember this is a frequency domain technique that operates on these switched signals and that means noise folding will take place and so you have to tell the simulator how far out should you grab noise and fold it back. Okay and that is what this does.

64. PSS Waveforms(Clocks)

The PSS then for example you have to look at the clock wave forms. This is part of the steady state output that would give you the timing of the circuit. Okay let me now grab the noise at this instant in time.

65. PNOISE Simulation Setup

And this is how you set that up. So in PNOISE you check the option time domain. This says take the sample of the noise and take it at that time and the importance here is I am taking the sample because I am not interested in the noise of the continuous time wave form. This is something that I indicated initially I want to know the noise at this point when the switch opens okay as a discrete time quantity because my signal is also in the discrete time. And again there is another parameter here that says how far out do you need to look now. In this simulation it is called “numsidebands” and that is again very important. I will show you to compute it.

66. How to Chose Parameters

So how do you choose these parameters? So we need to look far out enough to make sure that whatever noise you have that contributes significantly to the inband noise and goes this far out. How do you calculate it? So I mentioned it initially the integral converges about a factor of ten beyond the corner of the RC and so I can plug in what I know about the RC in fact this is my smallest RC from the switch. I combine this. So my maximum frequency to consider should be in the end this combined ten over π times N times f_S okay? So for this circuit example I have f_S 100 MHz. my N is 50 for the sampling switches and this is about 3 okay. So this is about 15 GHz I have to dial in. In other words I have to make the simulator look all the way up to 15 GHz to make sure all this noise is shuffled back properly.

67. How to Chose Parameters

The same then holds for the number of side band. It is the same thing. So why am I not setting these parameter I don’t know to 17 THz just to be sure? Well because my simulation would run all day. That’s the problem and this is the critical issue with PNOISE. So if you have a standard PNOISE type 2 it is really important to dial in enough side bands but not too many because you will be waiting too long and the most detrimental thing to do is to not dial in enough side bands and I have seen that many times because your noise performance will look stellar. You will think you are this God of a circuit designer low noise but all you did is you didn’t know how to simulate it and this is dangerous. Because in a normal simulation if you make a mistake setting up a transient simulation the results don’t get better, right. They get worse you get oscillations and what not and you will look into fixing it. In a noise simulation if you made a mistake your noise looks beautiful and you just go home. And so this is dangerous and so take this seriously if your noise looks to good you probably don’t know how to simulate it.

68. How to Chose Parameters

So luckily there are some tools now out there on the market by BDA that can actually handle an infinite number of sidebands at very high speeds it is a different algorithm that runs so if you had access at such a tool then you don’t have these issues anymore.

69. PNOISE Result(SpectreRF)

So this is the output then all set and done I dialed in the right number of side bands and here you see basically I get 2 outputs. One is the power spectral density and duh you know it is white you know why it is white we talked about it at length and at a second output I get at this time I looked at the sample. This is basically the variance so it is close to 400 µVrms. So this is my output.

70. Transient Noise Using BDA AFS

Transient noise is actually pretty easy to set up. In this context there is really only one parameter that I care about. One is essentially the maximum frequency that I need to inject into these noise sources that are plugged into all the transistors and resistors and again I know noise folds pretty much from 15 GHz all the way down significantly and so I set this parameter also to 15 GHz So this is about the same thing I need to dial in as my max AC frequency and so forth. So once I do this I can simulate a few thousand of these samples. So if I have a fast simulator like the one from BDA this actually runs really fast and using these thousand samples I can plug the spectrum and get the variance and so forth and this is actually pretty easy to do.

71. Comparison

And now I did all this and I did everything right here of course and this is kind of a good thing for you to do at home if you want to practice in the calculation. If you recall I had 474 µV. That noise gave me 485. I got 475 from PNOISE and the transient noise gave me 477. So to an engineer these are the same numbers right. And you can at least explain quantitatively what some of these discrepancies are. That noise has to be a little bit higher than the calculation, why? Because in the calculation I neglected the noise from the sampling switch in φ_2. So I can kind of can quantitatively explain. Here I have a little bit less because I don't have infinite side-bands. Here I have a little bit less just because of statistics I just have a thousand samples if I take more samples I can converge more accurately to the true noise which is probably around 480 or something. At least in this circuit.

72. Advantages of TRAN NOISE

So just to mention it you know I have been using PNOISE for most of my career but I am changing to transient noise because it is easy to set up and it runs very fast mostly because a lot of effort has been put into making transient simulators much better in the past and not as much effort has been put into making PNOISE better. So you do noise using the same powerful engines that everyone is using. So transient noise has improved a little bit faster than others. The other nice benefit is even if you have a non-periodic circuit you can investigate this you don't need to play manually with the noises and somewhat intuitively inspect the wave forms. So one thing to emphasize too if I give you now this 5th order sigma delta modulator it is actually pretty difficult to get PSS to converge in most simulators whereas transient tends to converge simply because it uses perhaps a more advanced engine.

73. Second-Order DS Modulator

So finally you know now all this stuff comes to fruition and we will now apply it to the real fabric. So let's discuss a second order sigma delta modulator. How would you go about finding the noise now in this circuit. So sigma delta modulator you have 2 interesting transfer functions that you care about from a quantization noise perspective signal transfer function, noise transfer function. Ian this second order example the signal transfer function is just a delay z to the minus 2 and the noise transfer function is a second order differencing. So you are basically high pass filtering the quantization noise. That is the main idea. So how do I go about figuring out the noise in this circuit?

74. Integrator

Well first thing I do is I look at my integrator and in most cases my integrator will look pretty much like this charge redistribution stage that we analyzed expect that I am not resetting this cap periodically using a switch here at the output. I am actually separating it from the sampling circuit periodically so that I can just keep accumulating charge right instead of just redistributing and then resetting it. So that is the only difference and so if you take this into account and this is not a big exercise. The input referred noise looks almost the same as before except that this term doesn't include C_f because it is disconnected at the time that we open the sample switch other than that this is the same. So I have known this input referred noise and one important comment to make is this input referred noise is white because the amplifier noise that is injected in φ_2 doesn't care about the history it just dumps extra charge and packets again and again on this cap. So I can refer all of these noises into the input. Referring them to the output would be tricky because I have an integration going on there. So working at the input is easier because there you know the noises are white.

75. Second-Order DS Modulator

So once you do this you know the following. So this is pretty high level and pictorial so basically you have this architecture for each integrator you compute its input referred noise variance and you know it is white. So this white noise for the first integrator is injected here. The white noise for the second integrator is injected here and now all you need to know is what is the noise transfer function from this point to the output. We happen to know this because for this noise source it is just the signal transfer function directly at the input. For this one it is a little more tricky but also not hard to understand the noise transfer function is basically a first order differencing because you have one integrator in the feedback path okay. So the first order differencing here. So I have white noise sources with these noise transfer functions

76. In-Band Noise

And so this is how I can think about it. At the output of the sigma delta modulator I see the first integrated noise basically unshaped. I just see the whole white spectrum. From the second integrator that one went through first order differencing that means it is high pass filter. So of course not exactly a triangle but I didn't want to draw this curve here. So it is basically noise shaped. These noises then hit my digital filter at the output of the sigma delta modulator and this filter takes away some of the noise. So from this white spectrum it just leaves me this slice so this is my signal band and so whatever this is depends on the oversampling ratio. And from here it just leaves me a little slice of this triangle.

77. In-Band Noise

So the combined noise, the overall in band noise of this sigma delta modulator is basically the area of this red rectangle and the green triangle and you can show that it corresponds to these expressions. This is just the total input referred noise somewhat proportional to kT on C times v_{in}^2 divided by the oversampling ratio and here because of this geometry of this noise shaping you actually get OSR squared and now you combine it and this is your total noise expression, okay? Little fine print that was not included here but you will figure out when you do this. The sample noise of the second integrator and the redistribution noise of the first integrator are somewhat coupled and so you need to make sure you are not double counting the two noise sources when you compute those white spectrum but that is a minor detail that you can easily manage.

78. Summary(1)

Alright this is the end I am already going over time I am sorry but what do I want you to remember? The first thing significance of kT on C noise steep trade-off with power dissipation. Don't fool around over designing

6dB or something like this. This is not a good idea. Remember kT on C noise is white in the spectral domain. There are 3 slides of math that show you that that must be the case. You never have to look at them again but again you can do it. Noise folding that brings this penalty of increased power spectral density and that is what people complain about most of the time in switch cup circuits. That is kind of an inherent penalty compared to continuous time circuits. Charge redistribution circuit analysis remembering that the information is in the charge domain and to assess the noise performance accurately you better assess the noise that is added to the sample charge and not the voltage. This can be detrimental if you do that wrong in the wrong circuit.

79. Summary(2)

The sigma delta modulator example you know it was pretty quick but I think you saw how once you have your fundamentals down these circuits can be analyzed in a nut shell on the back of an envelope especially when you work with input referred noises and you evaluate the transfer function. The same approach by the way works almost for arbitrary SC circuits I picked a somewhat more complicated when you can also do that with filters and so forth. In terms of simulations I showed you that there are at least 3 methods that allow you to simulate the stuff pretty accurately. Ss you saw there is a complete agreement between all methods including hand analysis but I would evaluate transient noise if you haven't done it for a while because these tools have improved significantly.

80. Further Reading

So finally this is not the end of your career in terms of noise if you want to know more this was just the intro tutorial. So Richard Schreier has written a nice paper that has most of the stuff in it that I talked about and in addition it doesn't make some of the approximations that I talked about. He has expressions that shows you what the noise from the switches is. He has some output referred noise components from the integrators and sigma delta and so forth and this is for cracks if you want to read this and get sort of the second order understanding. If you design more complicated OTA's then first order which I think you do you can look at this paper to find out what the kT on C expression looks like for those and finally this is a white paper by BDA that shows an example of a relatively complicated delta sigma modulator and how it was simulated using transient noise and it gives you an estimate on how long this takes and so forth. Just a pretty nice example to consider.

81. End

So that is the end of my presentation. Thank you.