ISSCC 2012 Tutorial Transcription
Managing Offset and Flicker Noise
Instructor: Axel Thomsen

---

1. Managing Offset and Flicker Noise

Thank you, welcome to the tutorial; I will be talking about managing offset and flicker noise.

2. Outline

So here's an outline of my talk. I will first spend some time talking about motivation and history, and talk about offset and flicker noise background information. I will then talk about three different techniques to manage offset and flicker noise, which are chopper stabilization, autozero technique, and a little bit about correlated double sampling and present a conclusion.

3. Motivation: High Gain Example

So let's first talk a little bit about motivation, why do we care about offset and flicker noise?

First example, we want to do something with very high gain. Imagine you have a 1 mV signal coming out of some sensor, 2 μV of noise, and you would like an amplifier with a gain of 1,000 so you get your signal up to 1 V, gain some noise up to 2 mV, so we'll use a 10-bit ADC and we'll be happy. But if we have a poor opamp with say, 9mV of offset, it limits the gain that we can implement, 900 mV of amplified offset, 100 mV of amplified signal, and now we've really made our problem harder; it's a 13-bit ADC problem, so the lack of offset would allow us to have a much simpler system.

4. The year 1956: HP 15-1800

So here's an example that I dug up from a good place to go when you want to look up the history of industrial, of measurement techniques, HP archives, and you see from about 56 years ago, exactly what I was talking about. We have a gain of 1000 amplifier and in order to use it, even back then, some sort of stabilizing amplifier stage with some sort of chopper so it was used in order to make this work.

5. Motivation: Small Signals, Low Noise

Another example: low noise, now let's think about maybe a Hall Sensor, produces μV of signal levels.

The flicker noise in a FET-based opamp, can have pretty large noise in a 0.1 to 10 Hz bandwidth, down here where the noise spectrum of the MOSFET shows a lot of flicker noise.

Thermal noise would be much, much smaller in the same band so the flicker noise would bury the signal of interest, it can be more than an order of magnitude above the signal level.

6. Motivation: Ratiometric Example

And then, maybe a third example, ratiometric measurement, very common in sensor interfaces. Imagine you have a bridge transducer where whatever property you're measuring, be it temperature,

pressure, or strain, is represented by variations in this resistance, the typical measurement set up uses the reference voltage here, and we measure some sort of divided version of the reference voltage at the input of the amplifier, and if we didn't have offset, we would simply get an output out of the system, that's the gain factor times the change in resistance, exactly what we wanted, but the presence of offsets, brings in not only these offsets, but also the reference and so the nice ratiometric measurement that we wanted, is basically ruined by the presence of offset.

7. The bipolar and JFET days

So, there was a time before MOS became prevalent where bipolar and the JFET devices dominated opamps, and back then, offset and flicker noise were not that big of a deal. If we look at the device back then, I would consider sort of the perfect opamp, the OP627, well it provides an offset that's less than 100 µV, very low offset drift, good noise density at even low frequencies. So noise and offset, really were minor problems at that time.

8. Comparison to a MOS opamp

So we come to the MOSFET technology, and I just picked some example MOSFET with typical performance parameters. Offset now is in the mV range, offset drift, which often is an even worse problems than offset itself has gone up, the noise density has gone up substantially. You can see here in the noise spectrum that the noise at low frequency is much higher and the corner where we change between thermal noise and 1/f noise, has moved significantly higher.

9. CMOS device characteristics

So, where do all these characteristics come from? Well if we look at the MOS device characteristics, we have to deal with a lot of different things. What we're really looking at is variation in the current flow.

There are surface charge traps that can capture some of this current and release it. We have to deal with the variations in the doping in source drain, bulk, gate, we have oxide thickness variations, all these manufacturing tolerances, they all get reflected in this threshold voltage of the MOSFET and that really leads to all our problems that we will see and that we're trying to solve.

10. Offset

So let's first talk about offset and how we quantify offset. So a very important paper in analyzing offset is by Pelgrom, shown in the references in the back.

The key idea is that really the threshold voltage varies by a factor of V_{OS1}, some constant, divided by the square root of the gate area. So V_{OS1} is a standard deviation of the threshold of a device with 1^2 µm gate area and these days, the data for that are given by the fab, and it is assumed to be based on a best-case layout.

Not only do we have offset, but there's offset drifts over temperature, so we cannot simply say we have this fixed offset, there's this variation, it varies with temperature and if you scan through a couple of data sheets, probably up to about 50% of the original offset could also be the offset drift.

On top or all this, there are things that I'm not going to talk about much which is, we can get additional mismatch due to layout effects, basically what we assume is that we have identical neighborhoods of the devices that we are looking at and when these neighborhoods are not identical, distances to well edges,

and active edges and the nearest poly and so on, will simply add to that.

11. Input referred offset example

This is kind of a busy slide showing how to do offset analysis and typically what you find in a lot of textbooks is 1/f flicker noise analysis. Offset analysis is very similar, but is often not given in textbooks, so I'm going to spend a little bit of time going over it.

So each device in the circuit here has a standard deviation of its threshold voltage proportional to this constant that I just talked about, divided by the gate area. So all these threshold mismatches, when multiplied with the transconductance of each MOS device, result in an output mismatch current. This output mismatch current, we can reference back to the input by dividing the value of this current by the transconductance of the input device. And that gets us to an equation that looks like this where the total offset is a combination of the offset of the four devices M1, 2, 3, and 4, where the contribution of device M3 and M4, is scaled by g_{m3}, the transconductance of M3 and g_{m1}, the transconductance of M1.

Very similar to noise analysis except that the parameters that we're talking about are offset voltages and we can plug in our equations and solve it and we get essentially an offset equation that depends on the gate area of the input devices, the offset constant here called A, and then the input-referred component of the lower devices M3 and M4 and we can simply plug in the numbers and we get a few mV of offset in this example. So, as I said, very similar to noise analysis just not often highlighted in books.

12. Flicker noise or 1/f noise

Flicker noise, what is flicker noise, well it's often also referred to as 1/f noise and a simple explanation is that it's a surface effect, that we have carrier recombination at surface traps, and so, that makes the current flow through the device have a 1/f type of a noise characteristic. So the very basic model that we find in the simple SPICE models is pretty coarse, does not really take into account bias dependency and such. It simply has two model parameters, KF and AF. AF is somewhere close to 1, and then the equation implemented is Vn, and this should really be Vn^2, is equal to KF/sqrt(W/L) divided by the frequency, to the power of AF.

And this is why flicker noise is called 1/f noise; it's proportional to 1 over the frequency, approximately, if this factor AF is equal to 1.

Now we talked about offset 1/f noise and really, when it comes to circuit techniques, they are one and the same thing. Offset really looks like static 1/f noise, whatever works on 1/f noise, works for offset as well. Any kind of technique that substantially reduces the 1/f noise, I kind of think of offset as sort of 1/f noise at DC, just stuck in time.

13. Noise spectrum of MOS device

And here is a noise spectrum of the MOS device. You can see here, at high frequencies we have the thermal noise floor, essentially given by the transconductance of the device. At low frequencies we have the flicker noise, which rises about 10 dB per decade, between 2 decades here by about a factor of square root of 10, and there's a term called the noise corner which is where thermal noise and flicker noise are about identical and we talk about this noise corner because it helps us understand, for bandwidths wider than the 1/f noise corner, really the thermal noise matters and for bandwidths lower than the 1/f noise corner, the flicker noise really matters.

14. Flicker noise analysis example

Flicker noise analysis example, looks very much like the offset example and is also covered pretty well in your average textbook. You can see here, that this equation looks exactly like what I had on the offset equation, except we're now talking about noise densities. We plug in our noise equation from up here, and we end up with a final equation that has the noise of the input devices plus an input-referred term of the load devices.

15. Flicker noise analysis example

And when we quantify it, we can solve this equation for the noise density at a specific frequency. You can see here, I did make the choice of a certain frequency and we have the noise density at that frequency.

16. How to manage offset and flicker noise

So now that we know what we're dealing with in terms of 1/f noise, in terms of offset, how do we manage, how do we build circuits that are precision circuits, in spite of these effects?

And I'll talk about a few things that we could do, and in the underlined parts, I will elaborate a little more about.

So basically the first three bullets have something to do with separating the frequency where the offset and flicker noise are present, versus the frequency where a signal of interest is present.

You can think about audio for example, DC has nothing to do with audio so you have a separation in frequency between the audio signal and the DC signal that you worry about, so a simple technique here would be AC coupling.

In communication systems, you might plan your system such that with a zero IF you bring your signal, maybe down to DC, which would be bad because it's where your offset and flicker noise dominate, a low 1/f implementation would still keep the signal away from your problem area.

And what I'm going to talk more about is chopper stabilization where we use an upmodulation of the signal above the flicker noise corner, in order to separate the signal of interest and the error effects.

The last three are shown here, are more of the type of technique that is, let me measure the offset, and then remove it. Right, we can do a one time offset calibration for example; we can do a periodic offset calibration and this is really what autozero is all about, and then we can separate the signal and error through correlated double sampling, so those are techniques where we actually sample the error and subtract it out and that's what I will talk about in the slides going forward.

17. Outline

So now this brings us to chopper stabilization.

18. Chopper Stabilization

And the key idea of chopper stabilization is illustrated here. We come to this in here, with an input signal, and what we're really interested in this input signal, is some low frequency signal. You can see

here in the input spectrum, we have some low frequency signal that we want to measure and amplify.

But in the same spectrum area is this red thing the offset and 1/f noise. So the basic idea is one of modulation. I have a modulation waveform here, that's a square wave, and I here, upmodulate the input signal, so that it ends up showing up here in the spectrum around my chop frequency away from the signal components of 1/f noise and offset shown here in red.

I do my amplification in this block, and I do a synchronous demodulation, or downmodulation of the upmodulated input, and an upmodulation of the offset, so I basically synchronously put these upmodulated components back together, and they end up being the original signal again. And the other that happens is that my spectrum of noise in offset, now gets upmodulated and shows up around the chop frequency in its harmonics, and I have to be a little careful, but I have definitely now, amplified my signal and separated it from the error components. I said I have to be careful because the offset and 1/f noise is still present, I just moved it to a different part of the spectrum. So when I go forward and do something with the signal, I better watch out that I don't do just the opposite of what I did here which is, for example if I sample this output, sampling would simply alias everything that's going on up here, that I really want to get rid of, back and mix it with the signal.

19. Transistor implementation

So how does this look like in a real circuit design? It looks kind of complicated with two modulators, but fortunately, these modulators simply are a set of four switches.

So let's first take a look up here where it's still sort of a block level diagram, here's my amplifier, and I basically have some switches that are turned on in phase 1 they give me a straight connection from the input and the output, and two switches that give me a cross connection here, basically invert the polarity of the signal. And I do exactly the same thing here on the output, and that's these basically eight switches are my modulator and demodulator. So actually it is pretty simple.

If we now say okay, let's take a look at a real transistor level diagram. Well we have our input signal, we have a set of transmission gates that connect the input signal to the input, and during phase 1, we go straight and doing phase 2, we basically invert. Then we go through a very standard, in this case, single-stage folded cascode opamp, and on the output, we demodulate again, synchronously, and that's essentially all it is to implement a chopper -stabilized amplifier.

20. Chopper amp in SC circuit

So imagine now, we use this chopper stabilized amplifier inside of a switched capacitor circuit, just to increase the level of complexity here, and then we'll take a look at what really goes on in this circuit. I need to put it inside in some closed-loop circuit to really see the dynamics of what the chopper does.

So what I do here is I have a standard switched capacitor integrator, implemented with the clock phases 1 and 2, and then I use a chopper-stabilized amplifier that's chopper stabilized with phases 3 and 4.

And there's a couple of observations we need to make here and they are, for one, my switching frequency of the switched capacitor circuit, is twice of my chopper frequency in order to go through one complete cycle of operating my opamp in a certain polarity and then following it up by an inverted cycle during the next step.

And then, I want to take a look at what really is going on here at this summing node. Usually when I

want to see whether my opamp is really doing what it's supposed to do, the place to look is the differential voltage between these two nodes here. And it reveals the following, that this summing node is never quite at zero, it is at zero on average, but in any moment in time, it's either at plus the offset voltage, or minus the offset voltage, just according to the polarity of phase 3 and 4.

So I have to be careful, through chopping, I have not turned my opamp into just a perfect offset free opamp, and the other thing is, that really a valid full integration of my input signal, I only get a valid output once I've gone through a full cycle of integrating once with a positive offset, and integrating once with a negative offset. So there are only certain times in my waveform where the opamp is really what I intended.

21. Waveform

This graph shows some simulated waveforms, kind of confirming what I just said; the summing node is never at zero offset, it jumps around between positive offset and negative offset.

This is my input signal and what I'm observing here is what my output signal does, and the output signal, as you can see here, has an integration of my offset, it integrates up, it integrates down. On average it stays in the right place and every completed chopping cycle, it will give me the correct output.

And then, as I now increase my input signal, I will start integrating up at my output, but I have to be careful that I always wait for a full cycle of phase 3 and phase 4 before I actually take a look at my output signal and believe what it gives me.

22. 2 stage opamp implementation

A few more implementation examples. This here is a 2-stage opamp, 2-stage opamp we will use typically when we have a resistive load to drive or when we need more gain, and this example here, is a fully differential completely symmetric circuit.

What we really care about is offset removal of the opamp 1^{st} stage. The 2^{nd} stage of the opamp, its noise will be suppressed by the gain of the 1^{st}-stage so we don't have to worry about it.

So we implement our chopper switches here, and we implement our chopper switches at the output of the 1^{st}-stage, the second stage remains without chopping.

So there are two things that I want to point out here, is that A) this is a very symmetric structure and chopping really works well, it works best when we have full symmetry, any asymmetry will corrupt this. The other thing is that, compensation capacitors, basically any kinds of capacitors that slow down the behavior of the circuit, best remain outside of the chopper loop.

Our idea here is to modulate up our signal so we actually have a fairly fast signal going through the opamp and we really want to maintain the signal to be fast and not slow it down because slow transient settling waveforms are one of the limiters of the precision of a chopper amplifier.

23. Waveforms (inverting amplifier)

Again here are some waveforms and we can see, this is implemented as an inverting amplifier, we have a summing node that jumps around as we apply an input signal, we see that our output signal, given here in this yellow curve, jumps around and approximates over time, this blue curve here, that's the ideal

output.

But at no point in time, do we really see the ideal output. So we've upmodulated the offset, it is still present.

24. Chopper stabilized 2 stage opamp

Let me point out one pitfall, or one limitation of chopping here with this single-ended 2-stage amplifier example.

So a 2-stage amplifier can also be chopper stabilized when it's not fully symmetric when we really have a differential to single-ended conversion present. What we do then is we switch around which node of these two nodes here, is the input to the current mirror and which one is the output to the next stage. Unfortunately, there's a fundamental asymmetry in here, because the voltages on these two nodes are not always the same, even with our best design efforts, if there is a resistive load here, like it would typically be, the current through this device, does not match the current through these devices.

And this kind of asymmetry, leads to systematic offset and the chopper simply cannot help.

25. Asymmetry issue – caution

This sort of over-dramatizes the effect, through asymmetry, you can see here, that the output voltage does not approximate the ideal output voltage anymore and the reason is that we have differences between the voltage on the current mirror and the voltage on the gate of the output stage. So asymmetry is really the enemy of a chopper-stabilized circuit.

26. Parasitic switched capacitors

Another pitfall of a chopper-stabilized circuit or, really it's not a pitfall it's just something to look out for, is the fact that we also, aside from building this nice modulator that up and down modules are input signal, we've really built a switch capacitor circuit. This capacitor C_{p1}, gets charged to the input voltage on V_{inp}, and then during the next cycle, gets connected to V_{inn}, and so it delivers whatever voltage was seen here, a charge package that is proportional to that voltage to the node over here, so it really acts as if we had a parasitic resistance between the two input nodes.

The same thing happens on the output, we have a parasitic switched capacitor connection between V_{out+} and V_{out-}, and this resistor creates an issue when we're interfacing with very high impedance circuits. This resistor really creates an issue when we're trying to generate a really high gain and now we've added a parallel output resistor on our amplifier that limits the gain available.

So these are unavoidable side effects and all we can do are minimize the size of these capacitors C_{p1} and C_{p2}. The really bad ones would be capacitors C_{p3} and C_{p4}. Those we can do a lot about through good layout, through good practice, any capacitance that goes from the outside of the chopper to the inside of the chopper, or here to the inside to the outside, is a bad thing. It will reduce the effectiveness of the chopper, it will demodulate things that we really don't want to be demodulated.

So C_{p1} and C_{p2} minimize but they are unavoidable, C_{p3} and C_{p4} really eliminate through dgood layout techniques.

27. Chop frequency considerations

How do we pick a chop frequency? Well, this is actually a fairly simple task, chop frequency needs to be above the noise corner in order to remove all 1/f noise, lower chop frequency could be chosen but our chopping gets to be inefficient, and then up here, when I look at the input-referred noise, well this is really the region where the opamp runs out of bandwidth, so we really need to operate in a region where the opamp bandwidth, where the opamp can quickly settle.

So what we really want to choose is a chop frequency that's as close to the 1/f noise corner as possible and this word here should be slower and not faster, okay this little correction here.

So our choice of chop frequency should be right here where we get into the thermal noise, but we don't chop too fast in order to avoid having too large parasitic resistances and in order to have, to spend too much time settling the chopping modulation. If we're going too fast, we spend most of the time settling and we really don't see the noise boost that we want.

28. Chop artifacts

So let me spend some time talking about chop artifacts. So as you have seen, all my waveforms that I've shown, when we look at the output, the output is not a nice DC signal, it's a DC signal that's sitting at the right average but it has a big square wave sitting on top of it, that's actually pretty painful to deal with from various points of view.

You can imagine if I had some offset voltage, typically I have low frequency signals, I'm dealing with sensor interfaces here, so my signal might drive me to have a certain slew rate and bandwidth requirement, while my chopper really has much tougher settling specs than my signal would have. I would have to build a much faster, much more high powered opamp, just to settle, just to be able to drive my output square wave here because I amplify my offset by a pretty large gain factor.

29. Waveforms of chop artifacts

So this is what this typically looks like. The summing node has the square wave on it but it's V_{OS}. My output node tries to approximate this signal here and it does it, but really my opamp needs to drive this very large square wave here and so it needs to be a high powered opamp with a lot of bias current to give me the slew rates and such.

So these chop artifacts are very painful to deal with.

30. Eliminate the artifacts

So the idea is, we can eliminate these artifacts and we can do this by using the following thinking.

Imagine this as being an opamp, so this circuit here is enclosed in a feedback loop so there will be feedback from here, back to the input.

Now, just like before, we have a chopper stage here, so this chopper stage really has zero offset so that's a good thing, we have the unmodulated offset present here, but now we place a filter here, so we remove the chop artifacts that we have generated and because this is a high gain path, we will now at the output of this, generate an offset correction that is just the opposite of the offset that this stage here presents to the input.

So because this DC stable high gain path really forces this voltage, this differential voltage to be zero, we'll generate whatever is necessary out here to compensate for the offset in this path and with this technique, essentially we have a DC or a low frequency correction signal here that cancels out this offset, but the chop artifacts are gone, they've disappeared into this filter and now we actually get the kind of behavior that's more equivalent to the perfect offset free opamp.

31. Removal of the chop artifact

So here's an example from an ISSCC paper from 2006 that illustrates this point. We have a chopper input stage, then we have a switched capacitor notch filter, and we have a fast path here and through this amplifier, we inject the offset correction and that way we get a chop artifact free chopper-stabilized opamp.

32. Switched Capacitor Notch Filter

Zoom in on that switch capacitor notch filter. Out here, we come out of the first chopper stage with chop artifacts and now you can see that the timing here is such that we have the two chopper phases, phase 1 and phase 2, and we have these two switched capacitor phases, phase 3 and phase 4, which are 90° out of phase with the chopper signal.

So during phase 3, on this capacitor here, we integrate half of phase 1, followed by half of phase 2, and because of symmetry, these cancel out and you can see here that this voltage C5 goes up and down, and then we connect to the output during the blue phase, and during the second half phase we look at half of phase 2 and half of phase 1 and we integrate onto this cap and then we switch to the output during the blue phase.

So this implements a very simple, very elegant notch filter that takes out the chop artifact very effectively and through that, our offset is removed and no upmodulated artifact is present. Only during the blue times, is the output of the switched capacitor circuit, connected to the stage that follows and you can see that this blue phase does not have any ripple on it.

33. Another implementation

This is a similar idea, a few years earlier, instead of using a switched capacitor notch filter, this circuit implements a gmC filter that's a low-pass filter and also additional attenuation, but the idea is essentially similar with enough attenuation here through low-pass filtering and through brute force attenuation, the chop artifacts are made to be insignificant.

34. Chopper implementation details

This is a circuit implementation detail showing a little bit more detailed work on how to implement the chopping with a little bit more consideration of transients, you can see here that on the input side in this work, the chopper is a voltage switching arrangement, on the output side, in this case, a current steering implementation was chosen, during phase 1, we steer the current coming out of this circuit into one capacitor and during phase 2, to a different capacitor, and then an additional phase 3 was introduced, whose purpose was to basically throw away some of this current during the times when transients are going on between during the switch-over time, between phase 1 and phase 2, so that these transients, which might not be symmetric and those are the times when the circuit is not behaving according to theory, this current would not get integrated on this capacitor.

So this is more in line with what a typical amplifier is doing, which is a voltage in, current out, type of behavior.

35. Results for the Cirrus CS3001

This circuit was implemented in the Cirrus logic CS3001 device, and you can see that with this attention to detail, low noise in band was achieved about 6 nV/rt(Hz), the flicker noise corner, through these techniques, can push down to about 0.05 Hz offset below 10 µV offset drift, less than 50 nV/^{0}C and I'd call the OP627, which was the opamp I like to use in the lab when I was doing a lot of lab work, the perfect opamp. Well, this kind of performance by far exceeds what bipolar and JFETs devices were doing. So it got to a similar noise level and by far, better offset drift and flicker noise performance.

36. Alternate proposal

There's another technique available to reduce the ripple that's present on the chopper circuits, and there have been a number of papers illustrating this technique.

If we look at the output here, signal and offset are still separate in frequency, so if we detect the presence of upmodulated offset here, we can use demodulation and integration to drive this ripple to zero.

So in a way, instead of filtering out the offset, we measure it through this demodulation path and drive it to zero through a closed-loop system.

37. Removal of chop artifacts

And this is a figure from last year's ISSCC where this kind of a ripple reduction loop was implemented in a current feedback opamp. You can see here, the chopper stage, there's a 2nd chopper stage here, here's the output where our ripple now would be present. We take this output, demodulate it one more time, now we can use integration to drive, basically an offset cancelling current into this stage, and the loop will guarantee that my output ripple will be driven to zero, or close to zero, by injection of an additional current into this stage.

38. Noise Analysis: PSS/PNOISE

So, these are techniques, another thing that we need to look at is can we verify it. What kind of tool is out to verify this?

And over the years, the periodic steady state analysis followed by PNOISE analysis, has really matured and this is a technique that should be applied to the simulation of all of these kinds of circuits, it captures all the modulation and sampling effects that are so difficult to capture in the typical small signal noise analysis that has been around for a long time.

So the idea is that the simulator finds a repetitive transient waveform, here it would be at the chop frequency period, noise analysis is performed with many timesteps along the way, and basically the transfer function and output noise are calculated, the sum of all these, and it is a really powerful tool that actually is very, very easy to use. The threshold of the learning curve isn't all that hard. It's very nicely integrated into today's tools and it really should be mandatory to run these kinds of circuits through this simulator.

39. Noise Analysis: PSS/PNOISE

So I have just some notes on what I set up for my analysis.

40. Noise simulation results

What I really want to point out is what we can see through this analysis.

So I have a noise plot just generated with a noise analysis and we can see the list of noise contributors in this circuit and they are all flicker noise contributors except for this one resistor here, and so that's really what I expect, this is the noise that I want to see.

And then we turn on chopping, and now I can't run a linear noise analysis, I run my periodic steady state analysis but essentially from the user interface, it doesn't look very different and I can see exactly what I had hoped to see, which are three things.

Chopping gives me no real thermal noise increase, the thermal noise level is about the same, the noise contributors shown here, are now all of the thermal noise variety, and my R48 actually has, which used to be just 2% of the overall noise, now is my top contributor at something like 60%, so you can see how effective this technique was in removing all the 1/f noise.

41. mixed signal implementation?

Last topic on chopping; is this suitable for mixed signal implementation? Yes, it can be done; demodulation can definitely happen in the digital domain.

The trouble is that we have a lot of bandlimiting effects, as we go from the amplifier through the ADC, and bandlimiting effects really make the chopping a lot less efficient, so this would lead to a choice of slow chop frequency because of these bandlimiting effects. So I think for an offset removal, this technique can be done in a mixed signal way where the modulation is done in the analog and demodulation in digital, but this is really not where this technique shines.

42. Summary

So in summary, it's a technique that has no noise penalty, I think the remaining transients that we see, all the switching back and forth, limit the accuracy. Chop artifacts are something to look out for, and switched capacitor connections between inputs and outputs, really are something to look out for.

And if you take this chopper amplifier output, and eventually sample it first, be very careful, you have not removed the offset, you have simply modulated it up.

43. Outline

So that gets us through the chopping and moves us on to the idea of the autozero, which is a fundamentally different approach to the same problem.

44. Definition of autozero

The definition of autozero really is a periodic recalibration of the circuit. So what we do is, during phase 1, we measure and store the offset. We really take the circuits offline to calibrate the offset. We store the result of the calibration, and then during a phase 2 which is the operating phase, we measure the signal and subtract the offset, and so now we have a calibrated circuit used for the measurement and

we have to provide a means to subtract the offset, analog or digital.

45. Autozero ADC example

So I think the easiest autozero example is presented with an ADC. Just simply imagine we have an amplifier and an ADC. During phase 1, we short the input to ground, we take a measurement of the offset in the digital domain, store it in an offset register, and then during phase 2, we connect to the signal and on the output, we subtract out the offset that we stored, very simple, this is the autozero technique implemented in digital.

46. Low gain comparator example

I'm going to present here, an example how to do it in analog in a low gain amplifier or circuit.

So a low gain comparator would be a typical example where this implementation would be used. So if we have an offset sitting here, we short the input together of our amplifier, we have a storage of the offset, of the amplified offset, A the gain times V_{os1}, we store it on these capacitors, and since these low gain circuits are often cascaded to give enough overall gain, here's a basically identical stage that follows it.

During the operation phase, we amplify the signal, we still amplify the offset, these capacitors hold basically the gain, times the offset, so they subtract that out and we get the gained up input here, repeat that one more time, and we get the doubly gained up input, minus the offset out of this technique.

This works fine for low gain amplifiers where the gain, times the offset is still a fairly small number.

47. high gain opamp example

For high gain amplifiers like opamps, we have to do something different, so what we would do here to implement the autozero is, we have a phase 1, where we short together the two inputs of the opamp, and we let the output serve to charge this capacitor to the output voltage.

Now that we have the offset voltage stored on this capacitor, we switch it over here in series with the actual offset, and now during the operation phase, this acts like an amplifier with zero offset.

48. Another opamp example

Here's a fourth example of how to implement autozero, and I'll go through the same thought process one more time.

During phase 1, the autozero phase, we short the input, in this case, we have an amplifier stage that produces the output voltage necessary to drive its offset to zero and then during the operation phase, we have stored the proper offset canceling voltage on this capacitor C_{az}, it injects the offset cancelling current here and we simply gain up the signal without any error.

49. Transistor level example

Here's a transistor level example of this kind of a circuit, you can see here that we have a 2-stage opamp, a telescopic cascode amplifier followed by a 2^{nd} stage. What we added to it is really a second

differential input pair and autozero switches and autozero capacitor.

50. autozero phase

During the autozero phase, we short the inputs to ground and we use the amplifier to basically find its own offset and store it on this C_{az} capacitor.

51. operation phase

During the operation phase, the offset is now stored on this capacitor and it injects the right current to cancel the offset that's present in the real circuit and therefore we get offset removal out of this circuit.

52. Waveforms

What does this look like in the time domain? Well, let's take a look at the various nodes, you can see here that we have one reference autozero node here, and then we have basically the difference between these two nodes, as the offset voltage that we're storing, in this case about 40 mV.

We have the input summing node you can see here that during the operating phase, my input summing node is truly sitting right at the differential voltage between this green curve and this purple curve, is close to zero.

During the time when we are in the autozero phase, I have a finite gain effect, an imperfect behavior of this circuit. You can also see that on the output, I have during the normal operating phase, close to ideal behavior, during the autozero phase, my output node drives to this voltage on the autozero node because that's what its job is during that autozero phase. And what becomes clear here, as my input signal starts to go away from zero and I'm in a normal operating behavior here, that half the time, my autozero opamp is actually doing what it's supposed to which is amplifying my signal. And the other half, it spends measuring its own offset.

So the key thing here, is that, yes the autozero works well, but I really only have my amplifier available half the time or a certain amount of time. It needs to be offline; it needs to basically go and get calibrated, and then come back and do my measurement.

So that definitely is something that needs to be taken into account.

53. Error sources

Now what are my big error sources in autozero? Well the biggest one is probably there's a noise penalty. So I measure my offset, and of course I do a pretty good job of measuring my offset but noise is present, noise gets sampled and so my offset measurement is not perfect, it is a noisy measurement. And typically if I just do what textbook autozero tells me to do I'll increase my noise voltage easily by 3 dB. So that's a pretty stiff penalty. To get 3 dB noise back, you'd have to spend a lot of power in your circuits to overcome that.

So this is definitely one of the big ones to watch out for when talking about autozero techniques. It has the potential to greatly increase the noise.

The second thing is, I have two operating phases, I have an autozero phase and I have an operating phase and I'm trying to cancel, I'm really trying to store what my offset would be during the operating

phase, so I have to be really careful that the differences that I have between the operating phase and the autozero phase are not too great. Different currents flow, switches are open, different glitches are present, so I really have to make sure that operating phase and autozero phase are as identical as I can make them. It can't be perfect, but that's a systematic error that might be present.

The third thing that I have to look out for is storage errors. Well if I have digital storage like in my first example, that's perfect, other than quantization error. Analog storage is a more tricky thing. I'm trying to trap a voltage on a capacitor and so all the sampling effects like clock feedthrough, like charge injection, like leakage from capacitors corrupt my storage.

So first line of defense in these kinds of techniques is typically a differential implementation, but it's definitely not a 100% problem solver because I'm going to store a signal that is not a zero differential voltage, it is a voltage that has a value so both sides cannot be identical.

54. Charge Injection

So let me talk about charge injection just for one slide, just to bring this problem home a little bit more.

So the channel charge here, depends on the gate to channel voltage that I have present here in my autozero switch, and the charge distribution of this channel charge depends on basically what kind of impedance it sees, so the impedance looking in this direction, versus the impedance looking in that direction is different and it might be signal dependent. So charge injection onto the capacitor is basically an additional error term on this voltage that I want to save here.

So if I'm not careful, I have a big transition here of a clock so that spans my whole supply voltage, so the amount of charge that's coming out of this channel is substantial and it can greatly modify the voltage present on this capacitor.

So definitely differential operation is recommended, and you can also implement some of these cancellation techniques as shown by this orange transistor here and I put a reference to this technique in the end.

It's not the world's most robust technique but it helps some to minimize the error of the analog voltage stored on this capacitor.

55. Charge Injection

So here's charge injection as viewed on the transistor level, right here is my autozero switch and capacitor where I want to store my voltage you can see here that I've made it differential. I've simply instead of straight applying my common mode reference on this side, I've also placed a switch and a capacitor here in order to have similar charge injection behavior over here so that in a differential sense, my error is not as large, but I will see voltage difference between here and here, so the amount of charge in this switch will be different from the amount of charge in that switch, and I will see impedance differences.

Here I'm looking into an amplifier output, over here I'm looking into some voltage reference that I'm looking at, so all of these will contribute something to the storage error so that my actual voltage stored on this capacitor C_{az}, is not exactly what I wanted to store.

So attention to detail in the sampling is probably the biggest, most important thing to focus on to reduce

the amount of offset present in an autozero system.

56. Noise analysis: PSS/PNOISE

Let's talk about noise analysis; it's a very similar set of slides as before, again autozero can be fully captured in the PSS/PNOISE analysis.

57. Results of PSS analysis

And I'm talking here about the settings that I have, but more importantly, let's take a look at the results.

You can see here when autozero is disabled, I'm flicker noise dominated here, and you can see in the spectrum that flicker noise is substantial. I turn autozero on and I see, for one, that the flicker noise is greatly reduced but also that my thermal noise is elevated, and that's a fundamental item when it comes to autozero.

58. Autozero considerations

So, I need a reset time for my autozero, so that's an issue. My circuits aren't available all the time to process my signal. I have additional circuitry to inject the offset cancellation, and I also need to watch out, I have a sampling process going on and when I sample, sometimes can pick up noise from places that I don't quite expect.

59. Autozero considerations

So let me talk about that a little bit first. So the example I've shown here was a 2-stage amplifier and if a 2-stage amplifier is poorly designed, there's actually a lot of noise present, due to the output stage of the amplifier. What we typically expect is that the input stage dominates the noise, but a poorly designed output stage could greatly dominate that.

And so when we sample the noise of multistage amplifiers onto an autozero cap, be careful. The PSS/PNOISE should highlight that, it's not a problem that can't be found with techniques present but it is something that is a possible mistake.

60. Reducing autozero noise

Next, let's talk a little bit about what can be done to reduce the autozero noise since that's a big limiter.

The key idea here is, that we could say that instead of simply sampling the offset and storing it, let's sample the offset and average it with the previous samples of the offset that we have and therefore, do something of a noise removal.

So instead of having a single measurement offset, we really do some sort of low-pass filter, here represented as digital low-pass filter, and with that kind of averaging, this offset that we're measuring here, gets to be a lot less noisy.

61. Analog implementation

We can do the same thing in analog, instead of simply sampling this offset here and storing it in one step, we now have basically an integration cap and a little switched capacitor that adds to this integration

cap so that this C_{az1} will simply hold a low-pass filtered version of my previous offset measurements, kind of a leaky integrator implemented in the analog.

And if I do that, sampling the offset and then adding a portion of that offset into this integrator, I can reduce the noise on that voltage very effectively, although I do in turn, reduce the bandwidth of my autozero and the effectiveness of the 1/f noise cancellation.

62. Reduced bandwidth and noise

So this basically is that effect written up in a few more words. Reduces the cancellation bandwidths but it also reduces the noise.

63. AZ in a continuous time opamp1

Next thing I want to talk about is two techniques how we can get over the fact that our amplifier is not available half of the time.

So here's a very basic technique where we say, hey, let's just use two amplifiers and use them in a ping pong fashion and we do the autozero on say, this guy here during phase 1, you can see I short him out and I do my autozero measurement here, and store my autozero voltage here, while the other guy is simply connected to the signal and is amplifying the signal, and then during phase 2, I take this guy offline, and I calibrate him, and this guy is doing the work, so I can kind of have a ping pong operation between these two amplifiers and I now have my autozeroed opamp available all the time.

64. Waveforms

Waveforms here, showing that this works pretty well, and we basically swing or jump around between, you can see the opamp here jumps around between doing the offset cancellation autozero operation, and being a perfect opamp in this circuit.

65. AZ in a continuous time opamp 2

Here's a different implementation, kind of a similar idea of using two stages but here the two stages aren't quite identical.

The first opamp here is working on the signal all the time, during phase 1, this amplifier A2 is being autozeroed and cancels its own offset and the measured offset of A1 from a previous measurement is stored here. During phase 2 we connect this amplifier, its offset now is cancelled so this is a perfect amplifier, it drives this node voltage to the right voltage to cancel the red offset and makes this amplifier, A1, look like a perfect opamp.

66. Performance data AZ-opamap

So this technique has been described in papers and been used in devices and you can see a good performance example here, offset of this device less than 1 µV, offset drift 2 nV/C noise floor 22nV/sqrtHz so a very good performance from the DC characteristics. Note though that the achievable noise floor with an autozero approach will not be near that of a chopper stabilized approach.

67. Autozero Summary

So autozero summary real quick, it's a sampled system, sampling has noise and aliasing issues and we deal with the fact that devices are taken offline to be calibrated, we need to deal with storage issues and analog storage has its challenges.

68. Outline

I'm going to briefly jump into correlated double sampling, but really, it's an extension of the autozero technique.

69. Correlated Double Sampling

Correlated double sampling is really applicable to switched capacitor sampling systems and it's an extension of autozero technique into the world of switched capacitors. The capacitor is already present so really the smarts are in how can I use these capacitors to store my offset voltages and then I can make one extension to the offset storage by saying, hey, I can take into account the finite gain of the opamp and also which can be considered a source of offset, and also can take that out.

70. Offset compensated SC amplifier

So I'm not going to spend a whole lot of time on this but for a matter of completeness it should be part of an offset presentation. So the idea here is that I already have capacitors present in this switched capacitor circuit and I'm going to use capacitors C_1 and C_2 as offset storage capacitors during phase 1 where this voltage now goes to the offset voltage and then during phase 2, the operation of the switched capacitor circuit, I subtract out the offset voltage of the opamp because I have precharged this middle node here to be already at the offset voltage of the opamp.

So it's essentially an autozero technique where we just take advantage of the fact that C_1 and C_2 are already there.

71. Waveforms

And here's an example that shows how during phase 1, the difference, the offset is stored on this summing node voltage and then it simply shows up shifted up during the second phase.

72. CDS for gain error correction

I can extend this and put a gain error correction into this. Essentially I add a capacitor C_3, which measures or which stores, when this output voltage here is sitting at V_{out}, it does measure what the summing node voltage, due to that limited gain, basically the voltage V_{out} divided by the gain of the opamp is present at V_A, and I store this on capacitor C_3 and during the next phase, I can subtract this voltage out and now my opamp really acts as if we have twice the gain.

So that's the additional capability of the correlated double sampling technique that we can also take out the offset due to finite gain.

73. Error sources in CDS

Error sources in correlated double sampling, again we have charge injection, we have sometimes noise

penalty, noise penalty is fairly small because we're already in a sampled system, only if we use additional capacitors we get some additional noise penalty here.

74. CDS summary

Really correlated double sampling is an extension of autozero with the additional capability to subtract gain error related offset.

75. Comparison

We get to this comparison slide, so chopper stabilization is a modulation technique, I guess that's the big difference between the chopper stabilization and autozero. We're talking about a modulation technique, no sampling is present, versus we have a sampled system, so we have no noise penalty in the chopper case, increased noise in the autozero case.

Accuracy, I would say the accuracy in terms of offset removal, I couldn't say the chopper does better, or the autozero does better, it's just the mechanisms are different. In one case, limited by transients, and probably if you survey the state-of-the-art, I think you do find the best offsets lately in chopper stabilized circuits, but its more about where do I focus on in order to achieve this accuracy. It's the transients in the chopper that really make the difference, whereas the precision of the sampling process is the key here.

If I had to make a choice, I think if I'm dealing with a continuous time application, chopping is the way to go. As soon as I add sampling to the system, chopping gets a little bit more complicated because sampling aliases and sampling is really something that's more natural to the autozero world.

So, in my observation, I would say if I'm dealing with a continuous time application, steer towards the chopper if you're dealing with a sampled application that really lends itself to the autozero technique.

76. Conclusion

So in conclusion, good low frequency and DC performance has a lot of value in applications such as sensor interfaces, and although MOS devices have poor matching and flicker noise, there are several techniques available to improve this.

The noise power and area penalties, they're there, but they are manageable and well understood. I think a very important point is that this is a fully simulatable technique so the modern simulation tools fully support this and allow full verification and the resulting noise and offset performance can be world class. I mean that the best offset and 1/f noise performance does come from MOS circuits, and that concludes my tutorial.

77. Papers to see

I have three papers to point out that are in this conference that are interesting chopper papers, so they're all in one session, all in a row so it makes it easy for you, and

78. References

I have a list of references which are some of the papers that I mentioned as background work for my tutorial, others are just state-of-the-art DC stabilized circuits that are of interest.