ISSCC 2011 Tutorial Transcription
Interfacing Silicon with the Human Body: A Primer on Applications, Interface Circuits and Technologies for the Medical Market
Instructor: Tim Denison

---

1. Introduction

2. ACKNOWLEDGEMENTS FIRST!

So, I want to start with acknowledgements first because it always gets short-shrift at the end. I have some colleagues that I really want to bring out who have helped me actually create this tutorial. It's a sum of all of our work together.

In particular, Chris Van Hoff and Firat Yazicioglu at IMEC gave a significant slide and thought contributions to this work. Then I pulled on the work of Pedram Mohseni on Chemical Sensing in the Brain; Albrecht Rothermel and Professor Ortmanns on Actuation and Implants; Professor Sarpeshkar from MIT and gave me some thoughts and the slides on low power circuits; as did Joyce Kwong and Anantha Chandrakasan on Digital Signal Processing. So, I want to start out by thanking them for their assistance.

3. Interfacing "Electrical Signals" to the body

So I want to start out the day, because it also, I won't say it's an apology, but, we're joining a time-honored tradition of interfacing electrical signals to the body. So back even in Roman times for the treatment of disease, they would prescribe things such as standing on an electric eel to treat your gout. And, so, this idea of actually putting signals into the body and treatment of disease disorders has been around for a long, long time. So, it actually makes it challenging to put together a tutorial.

The first version of this talk had about 300 slides. So I blame Willie Sanson, you know, he's very popular right now on the cover of the Journal, so, we can blame Willie for cutting back. So, it's quite difficult to try to capture all the challenges of interfacing circuitry to the body in one hour-and-a-half tutorial. So, please feel free to email me with any questions as follow-up, and I'll do my best to cover those. I'm going to stay around for an hour after the talk as well.

4. Interfacing Silicon to the Body, A Systems Approach; Tutorial and Architectural Overview

So given that challenge of how do we go about defining and trying to understand the constraints of interfacing the nervous system, I thought it was kind of good to draw a block diagram and a sensing approach.

Here's what we're interested in, is the body, and both detecting signals, and those interfaces or the signals that are going to come out; thinking about the sensors, the different sensors and interfaces involved; and how we process those signals. But, we're not just interested in pulling signals out of the body; we also want to have an effect on the body. We want to get about, going about the treatment of disease, and so thinking about what are the mechanisms by which to actuate activity within the body and try to provide some therapeutic benefit. I apologize. I just went through this slide in another practice; you can't see red on blue very well in the back.

We have to also think about the constraints. So, we want to sense signals, process them, close the feedback loop and control. But, we also have to consider things like the power that's involved; the kind of data storage retention; how we're going to transmit this data, both out of the device, or back into it for programs; a full

system consideration.

So, throughout the talk, we're going to keep coming back to three themes. So when you're a system designer in this area of implantation, or even an external device, there's three themes that always come back:

How are you going to manage the energy of the device?

How are you going to manage the information?

How are you going to manage the interfaces?

So we're going to keep coming back in all these contexts to those three topics.

5. "Actuation" Systems and Interfaces

But, to kick it off, we're going to focus, kind of, we have to focus a little bit in ninety minutes, so we're going to cover three major themes. The first, Chapter 1, is actuation systems interfaces. Then, we're going to talk about some sensor interfaces, and then, some of the signal processing techniques that are involved. And then, at the end, I'm going to give a couple of examples from industry right now of how these are put together.

So the theme is to try to cover these, kind of, give you a broad brush stroke with a few deep dives that cover the topics that are very high interest, especially at this year's ISSCC. I'll touch on those as we go through the talk.

6. The Power of "Actuation"

So the power of "actuation", I kind of, I lead off with "actuation". Everyone always thinks let's start off with sensing, that's kind of cool, but in reality, this is a typical patient (I don't need any volume, thanks.)

This is a typical patient for central tremor. You can see within his body's feedback loop system, he has problems controlling his pen. And, actually, the more control he tries to bring to bear, the more difficulty he has in writing. So you can think about this like there's almost not enough phase margin in his control loop. We need to do something to help out his motor control system.

So, actually, an electrical stimulation can help address this circuit malfunction. And the reason I start out with a picture of patient is, one, that's what this is about, it's not ultimately about technology, it's about trying to serve our patients; but, the other thing is, this gentleman knows he has a problem. So, I don't necessarily need a diagnostic that tells him more precisely that he has a problem. The number one goal of the patient is to say "Fix my problem." So, I want to start the tutorial with a bias toward "Actuation" in the nervous system and in the body to fix problems.

7. Presenting Problem: Neuromodulation Design

So, here's the presenting issue. This is an example from a neuromodulation design. So I actually, this is a slightly different figure than from what's in your book so this makes it a little clearer on the screen. So what are the different elements of a typical brain implant today? Well, where the rubber meets the road, or where the electron hits the neuron, is the electrode system. And those can be either subdermal electrodes on most implants today. There are a few companies with investigation devices looking at surface electrodes. These electrodes then need to pass through a lead body and extensions, and through an electronic generator that does the job of storing the energy in a battery and has the electronic system that gives you the right pulses that come out and go to the tissue.

And, what's nice about this system is that it actually embodies the three high level topics we're going to talk about: What are the energy sources and how are they managed? What's the information, basically, what's the biophysics and the coding that we want to use to drive the network? And what are the interfaces we have to consider in designing such a system? We're going to kind of hit on those themes, especially with the circuit, from a circuit standpoint over the next few slides.

8. Deriving Specs 1: Patterns for Neural Coding

So, the first thing when you're designing this kind of system is what are your specifications? You know we're engineers – how do we come up with specs? In reality, at least from my point of view, the first thing is to understand what are the patterns in neural coding? So, this leans heavily on the biology. So, we need to understand what are the biological drivers behind a design, and merge those biological, anatomical, and physiological constraints, and marry them with an engineering, design and architecture. Oftentimes, it's the interface circuit that actually is the joiner between these two systems.

So, this is an example that I pulled from a cochlear implant. You can see the idea here is sound comes in. There's a deficit within the nervous system of the patient. We need to actually over-drive, if you will, the neurons within the body. We make an observation that within the cochlea, there's a frequency mapping. And what we need to provide is an electro-drive that's essentially the Fourier transform of the sound that's coming in.

So that thing guides us in the high level design of the system in terms of understanding the signals, the way we need to process them, filter them, drive them into the body through the electrode. One of the key constraints we have, though, is how do you actually take these high level principles of signal processing and bandpass filtering, but then, actually drive the tissue itself? How do you actuate the nervous system and that's given by this block here.

9. Deriving Specs 2: Waveform Parameters for Modulating Activity

What are those specifications that you need to consider? This is a typical pulse pattern. This is the charge flow, the current versus time, on the horizontal axis. The idea here is, you can see, generally, we have a bi-phasic pattern so we deliver a certain amount of current during one phase, drive it out of the tissue, and then, we have compensatory pulse, and a recovery phase.

This pattern has a certain repetition rate that is a function of the neural coding that you need to provide, and is driven, oftentimes, by those higher level specifications. But, oftentimes, a circuit designer says are these arbitrary? How did we actually come up with these pulse waves and these different shapes?

So, at a high level, we have the bi-phasic nature because we need to ensure charge mounts for electrodes. But the actual details of what's going on within the pulse are derived by the electrophysiology.

10. Primer: Physiology of Electrical Stimulation

So, here's your one slide primer on electrical stimulation. If we step back and kind of look at the operation in the nervous system. It is an electrical system. It's set up by a gradient of carriers, if we will. Sodium and potassium are the primary ions. There are other ones that come to bear like calcium, magnesium, but sodium and potassium are good enough to get started with.

The gradient of these ions across the membrane look very much, and set up very much, just like a PN junction, if you will. It's kind of a stored nervous potential and that is a battery. It holds the capacity across the lipid bio-layer and that is basically a stored up reservoir of energy that we can tap into. The way we do

that is by kicking off through a stimulation electrode an electrical potential. What that tends to do is de-polarize the inside of the neuron, kick off a chain reaction through the sodium channels, and that initiates an action potential. And, so, the key toward driving and actuating a nervous system is to provide a stimulant that basically opens up the sodium channels, and then, kicks off the chain reaction and initiates an action potential.

11. Example: Neural Activation Functions

The actual duration that you have to provide of the current to achieve that channel opening and the amplitude that's required, can kind of form a relationship. You can kind of get some intuition on this. This is the idea of the classical strength duration curve. This is derived from Rothermel's paper in 2008 where he went into a retinal implant and then started to measure activity within the neurons and said "What is the amplitude, what are the constraints on the amplitude and the duration of an impulse that I have to provide in order to fire off activity within the neuron cell itself?"

This gives you the classical, kind of a parabolic, hyperbolic curve. The stimulus versus, the stimulus strength versus the stimulus duration and you kind of get the "sweet spot", the chronaxie, which is a point of minimal energy to drive that actuation.

So these are actual recorded responses from the retinal cells with Rothermel's implant and you can see the chronaxie going down along this axis. So, it's that physiology that responds to the neurons which actually sets up the constraints for the amplitudes and the pulse widths that one has to provide. So, as a circuit designer, you need to get in and talk to your physiologist and understand the constraints on your system so that you can design it appropriately.

12. Typical Stimulation Requirements for Neuro and Cardiac Stimulators

So, here's some of the typical stimulation requirements that you see within a neuro or cardiac stimulator. Kind of gives some of the degrees of freedom, well, the similarities and the differences.

The classical cardiac pulse generator, you know, you are all familiar with your heart rate, so it probably isn't too much of a surprise. They beat between 0.5 Hz, 3.5 Hz. This is, kind of, the frequency at which you provide the stimulation. The current amplitude, 1 to 10, 15 mA. And with the duration of the, on the order of 100 μs to a ms. For the pulsewidth, they tend to be on the order of 2, sometimes 3, electrodes. So, they'll drive, here, your atrium and your ventricle together on the right side of your heart. And, then, the amplitudes that you want to end up sensing in the presence of that situation that you want to end up sensing on the order of hundreds of μV larger.

Now, what's interesting is, the actual background current for stimulation is the stake in the ground. So when you start thinking about energy, what is it that sets "what's low power?" You go around ISSCC and you go to twenty different sessions and you hear "low power" in one area, its low power is 10 mW, and in another session, low power means a μW. So, what defines low power? That's a big open question.

From my point of view, it's what does it take to actually drive the physiology. So, in the case of a cardiac implant, 10 μW is kind of your stake in the ground. That's what it's taking in order to drive the system. But, when you look over at a neurostimulator, you see that the frequencies there are much higher. So, we're driving the tissue at a much faster rate, but, with kind of the equivalent stimulation amplitude. So what that translates down into is that in a neurostimulator, often the powers are orders of magnitudes higher, on the order of 100 μW to a mW of energy. In fact, some of the systems that are being explored for prosthesis are even 10X that, on the order of 10 mW.

So, when you're thinking intuitively about what sets the power, it's derived by that physiology and the need to drive the tissue and that gives you a metric for how to design a system and kind of, when you're doing your back-of-the-envelope equations, can you use an off-the-shelf amplifier or are you going to have to go custom?

13. Interfaces: How Do We Activate Tissue?

OK. So, this is a very abstract. We've got the physiology. We understand kind of the rough powers. Now, how do we actually make the connection to the tissue? The trick is, tissue conduction, so going in either the heart or the nervous system, is ionic in nature; and electronics is driven by electrons. And, so, we kind of have a little bit of a mismatch. So how do we take in our electrical system, using electronics, electrons and holes that then derive that into a system that's actually talking in an ionic space?

There's two methods to go about doing that. The predominant one is shown on the left hand side. This is to use a material, a noble material, that's non-reactive and basically polarize it. So you can think of this as a double-layer capacitor. You get your surface hydration, and then, through a double-layer capacitance, here's your electrical metal interface, you basically draw and repel ions within the solution. The idea here is to make it as non-reactive as possible so that you don't get any long term corrosion. So, this is the current electronics, that are in deep range spinal cord stimulator, are using platinum as their base with a little bit, 10%, of iridium, but that iridium is for mechanical, so, kind of, platinum, stainless steel, had been the workhorses for these kind of electrodes.

Another option, which is attractive, but is not used quite as widely for any kind of chronic stimulation, is to actually do a non-polarized system where you're getting a reaction at the electrodes. You actually have in this case, the classic silver-silver chloride. There's a chemical reaction that basically does the job of translating, transferring an electronic conduction into an ionic conduction. But, that's not quite, doesn't hold up quite as well in a chronic environment.

So, who cares? I mean what, why does a circuit designer care about electrodes? So, the big ones, especially in using the noble materials, is that when you're trying to drive tissue, polarization creates a big over-potential that you have to drive over. You have limitations in the voltage headroom. Like any capacitor, you can only drive so much charge through it before you get into problems. And, you need to worry about the safety of that charge transfer because you can actually kick off unintended reactions. Even when it's quote, "a noble material", so you have to be a little careful.

14. Some interface Constraints: Multiple Standards

So, that gets into a constraint. So as we go through the tutorial, every once in a while, I'm going to throw in some constraints that you should be aware of because it'll actually help you judge the quality of some of the ISSCC papers, and it'll also, that's not really, of course, your number one priority in life, or it shouldn't be. The number one priority should be, also, you're designing a safe system for use.

And, so, one of the key constraints when you're deriving a system with these electrodes is to understand what the charge density limits are that you can provide. There are two methods that are used. One is just the absolute charge density per phase. This has a long history with the regulatory bodies. It favors larger electrode systems. Another approach that's being used is this parameter called K, which is then derived by both a combination of the net charge per phase and μCoulombs, as well as the charge density, and looking at the summation of those two equaling a constant. That is another variable that we keep track of when we're designing these systems and use it to compare to predicate devices that are out there.

One of the challenges, in terms of kind of giving you a broader feel for where we're at in the design of these systems, is that a lot of the electrode safety is derived from historical experiments that were limited to

relatively acute times. And, so, because of that, a lot of work is designed and justified based on predicate work. So, for those of you who are in an academic realm, an interesting area for exploration is to actually to go and have a better understanding of what are the ultimate charge limits of different materials and their chronic performance in the body.

15. McCreery: intracortical and cortical stimulation of cat parietal cortex safe to K=1.7

Here's a publication that came out relatively recently, in 2005, where they were going through and looking at summation of, basically, and justification, of all publications and different studies of electrodes, and trying to define what is a safe K factor for chronic electrodes. So for those of you who are interested in getting a deeper, much deeper understanding, I'd point you to this reference and the *Journal of Neuroscience Methods* .

16. System-Level: Generating Stimulation Signal

OK. So, we've got high level specifications. We've got an electrode system. I've given you some of the constraints to think about. Now let's talk about how do we actually design the circuit; how do we actually drive the tissue. So, here's the picture. (OK, it's in your book. It looks a little fuzzy on the screen.) Here's our bi-phasic pulse and I've got them in two phases where I've labeled red and green. The red is stimulation of the tissue. Green is the recovery to get the charge balance from those electrodes. We need to provide energy into the system, flowing into the IC. I put this charge pump, and we'll talk a little bit about that, because oftentimes the voltage levels of the energy source are not high enough to actually drive the electrode system in the tissue. So you need some mechanism in order to actually boost the amplitude up to drive the network.

Given the two phases, at the very highest level you tend to have a current source driving in one direction through the tissue, and then driving down to a return electrode. It's not shown here for the sake of simplicity. These, in some modern implants, especially from the company, Boston Scientific, their advanced bionic student division, they actually can split this up into multiple different electrodes. So, you'd say "Why do I have a sink and source current?" Well, that's because it's not quite that simple. They actually split up in across multiple electrodes and that allows them to steer the field and actually have quite a bit of control. That derives from their cochlear implants as well.

So, then on the rebound, we recharge. Basically try to drive an equivalent amount of charge back through the system to keep the balance, the net charge balance between the electrodes, in constant. That is called the bipolar mode, so I'm putting it between two electrodes. If you are only going to return to one electrode, you don't need to go through the complexity. CASE is kind of our big, that's the IPG electrode CASE, that's a nice big return electrode. You can just throw a switch there. That's called unipolar mode. That's sometimes used in implants where you're really worried about efficiency because you don't have the overhead of another current source.

So, to kind of set the stage, before we talk a little about the transistors, there are a lot of capacitors floating around this system. Right now, there's a capacitor on every electrode. The typical neurostimulator today has 16 electrodes. They're looking at getting those up to 32, 64 electrode systems. Before you know it, 80% of the electrical content on a circuit design could be a capacitor. So that's something I want us to think about going forward – how do we get rid of those capacitors? They're a nice way to ensure net charge balance so that we can prevent any corrosion or tissue damage from phase to phase, but one of the big challenges for integrated circuit designers, and one of the true things that they can bring to bear, is to eliminate those capacitors.

17. Example of a Stimulation circuit ("Stim Engine")

So here's a stimulation circuit. We call them "stim engines" in the field for a, that was pulled from last year's ISSCC paper. They're pretty straightforward. You start out with a reference current for the system which kind of sets your master reference. And, then, the key is to drive that current in a controllable manner through an output electrode array. The reason that I took the effort to actually show this level of detail is that we're exploring this kind of current source where in phase C1, we kind of look at the reference resistor, sample onto a capacitor, C, onto this transistor, this PFET here. Then during phase zero, we actually deliver, and you can think of this as a little servo loop, we deliver the voltage and servo that was across the reference, drive that across R_0 so that we basically get a referenced current ratio between those two resistors.

So the advantage of this circuit is that, kind of in that sampling phase, we do a little bit of auto-zeroing of the transistors within the network which is good in terms of our scaling and getting our maximum accuracy. The other nice thing about the servo loop is that, unlike just a simple current mirror, is the benefit of the feedback. We actually get much better compliance. So, there's much less of a headroom requirement over the output current source than would be prior to if it was just a simple current mirror.

So, those are the kind of things when you're designing the actual output drive you need to think about. But, can avoid, what's the accuracy requirement? How much do the sinks and sources need to match? And, then, how much compliance do I need? Because extra compliance is power. You don't want to actually be burning 30% of your current within a device in the transistor circuits. You want every coulomb, the ideal is that every coulomb that comes out of the battery goes into the tissue. That's the kind of system that you want to derive.

18. Conversion from Energy Source to Electrode Drive

Now, there's a trick. The problem is, you know, we've got a stimulation circuit that may need to go from a few volts up to, you know, modern devices actually can go up to 10.5 volts for high energy discharge and some neurocircuits. Of course, the batteries aren't always supporting that. Typical cell voltage here from rechargeable, from, this is Quallion, where they made their characteristics public, you can see we go from 4 volts down to 3 volts, and then have a very steep discharge curve. You may be wondering what are these different "C"s that you see. This is actually the recharge rate. You can see the more aggressively we try to recharge the battery, so 2 C is basically double, think of the amp hour, if we're going to cut that value in half, that actually has an impact.

But, as circuit designers we have a challenge because the energy source that's been provided to us doesn't exactly match the overhead that we need to provide to drive the tissue.

19. Methods of Boosting the Voltage (all used)

So, how do you do that boosting? This is classic boost power supply considerations. All of these have been used in implantable devices. Capacitor stackings, so like all things, it's a design so there's tradeoffs, pros and cons. Capacitor stacks are simple. The disadvantage, the biggest ones we see, are you need very low impedance switches. And, you can think about it, when you need the most energy, so you need the highest rail, you're also stacking the most number of capacitors together so your effective capacitance is quite diminished. And, so, you end up having enormous capacitors within a device. So, that's falling a bit out of favor.

Inductively based, kind of switch mode power supplies, awesome. They had their time. They were very efficient, had ferrite core inductors inside. The problem is that they take a bit of space, and the other thing, though, is you need to think about your total system. A lot of these devices, the trend is both MRI safety, as a must, but, also, MRI compliance. And, so, you need to think about how will your system actually operate in such an environment? The inductive systems, with certainly any ferrites, are definitely falling out of favor

given those system constraints.

Charge pumps, the advantage is high efficiency. You can have kind of one large whole capacitor and then a bunch of smaller pump capacitors and adjust the rate accordingly for the energy that needs to flow. The biggest disadvantage, as least that we've found, is modest circuit complexity which I just demonstrate on this next slide.

20. X1 to X5 Charge Pump Circuit, and Example

This is kind of a simplified version of a charge pump that we have operating in one of our prototypes. You can see the different charge phase, pump phase, there's not a lot of circuit "know-how" that I need to pass along here. It's more just showing that you've got to think about all the complexities of the switch drives and the controls for the different methods. And, also make sure that you can actually drive those switches. You know, the classical how do I turn on the switch in the highest setting and make sure that it's operating appropriately?

We do find, you'll see some papers have tried, in some cases, tried to do on-chip capacitors. That's not typically done for anything that's moving energy around in our circuit because the parasitic capacitance is on an on-chip capacitor, the sizes just don't play well for efficiency and energy requirements.

21. When you get it right … Essential Tremor Patient Before & After (series of video clips)

So, now it's a nice point to pause. Remember why we're doing all this? And, here's, you get it all right and here's our patient again on the left. His stimulator's been turned off, and on the right, this is the same session, same day, his stimulator's been turned on. Through stimulation and the VIM of his thalamus, he's actually able to regain, not complete control, but much better control of his motor capabilities. That's what's sort of rewarding about working through all these requirements and getting it right is that you actually do have a meaningful impact on these patients.

The stimulators that I've talked about so far are pretty straightforward. The constraints are fairly well understood.

22. Emergent Application: Retinal Prosthesis, Advanced Stimulation/Actuation System

The next area for stimulation is getting some more advanced work. So, these are looking at an emergent application, specifically the retinal prosthesis. The idea of the retinal prosthesis is to treat macular degeneration, retinitis pigmentosa. The idea here is that the transduction cells basically within your eye have stopped functioning and so, instead of us relying on light to transduce the neurons and then drive electrical activity, we're going to actually use a photoreceptor of some sort, and we'll do the job of doing the transduction and the neural coding.

And, so, back to our three friends of the energy, information and interfacing. We're going to do spatial mapping, 1:1, we'll detect the photon and drive it. We need to do spatial derivatives because our eyes are functioning in terms of differences. We want to use logarithmic intensities and, of course, we need to actuate safely, slip it under your retina.

23. Concept for Retinal Implant

So, this is the idea that Albrecht shared in 2008 at ISSCC. He slips the electrode under the retina, and then, through a series of photodiodes in an array, can detect light, measure it's intensity and then use that to drive a titanium nitrate electrode. So titanium nitride, in terms of, from a circuit design perspective, it's quite a nice

material in terms of its polarization and charge drive characteristics. And, it's the de-facto material now in pretty much all cardiac stimulators. So, this is what the array looks like, and putting that under the array. So the characteristics are kind of two things I wanted to pull out of this work and drive home that are of interest for a circuit designer within a systems context.

24. Retinal Implant Supply Architecture

And the first is the implant supply architecture that they used. So when you're packaging a device, one of the big issues that you can run into is DC biases across material because those can lead to corrosion. So, just like you see on your motorboat or things like that, you'll have the little piece of metal that's actually your sacrificial corrosive metal. That's because the different materials in touch with a conducting solution have a DC bias build across it. So that's what we want to avoid.

And, so, within the circuit itself, what Albrecht's team did was, instead of having a DC bias across any of the external electrodes that are going to see tissue, they actually made it into an AC supply so that there's no net DC potential that's exposed to any fluids. Then they took advantage of that in deriving their output circuit. So that's something else for you to think about. It's not wise to have a large DC potential going through a wire in the body. It'll leave you susceptible to corrosion.

25. DC to AC Converter → Apply for Stimulation

So, within their stimulation circuit, as they discussed, they have, there are two things to consider. One is the neural coding, which is quite straightforward, you have analog input that's driving a differential transconductor. The nature of their photodiode, with the simple junction physics, gave them the logarithmic characteristics that your eye naturally codes for. Your level of brightness is not a linear function. We respond in a logarithmic manner. So, they get that for free, basically, from the junction physics and then they use that to drive a current in two phases. And, so, by switching the polarity of their externally exposed electrodes, they can do a bi-phasic switch which is shown in the book.

26. DC to AC Converter → Apply for Stimulation

So, just kind of flip it over, use the same transconductor and then you can drive through the reverse direction, so, pretty slick little circuit.

27. But No room for Capacitors! Solving the Charge Balance Problem

I wanted to give it a little more context in the corrosion. We talked a little bit earlier about that first design in a system for capacitor constraints, how do we solve that problem of capacitors? This is a big issue, so if you want to make an impact in this field, help remove capacitors from stimulation chips that are implanted.

So you can model your devices in one of two ways. The first, number one, is suppose you have an ideal electrode that looks like a capacitor. It's a bit more straightforward. You can just throw a switch and do a discharge. That kind of that gives you some level of charge balance. Historically, cochlear implants actually use this quite a lot between their stimulating gain firing that the electrodes that are exposed within the cochlea, between phases they would then discharge their charge, or discharge their electrodes together to ground.

The other kind of electrodes, and this can be more representative, is when you do have resistive path to ground, you need to get charge balance a little more carefully by keeping the net potential, I'm sorry, the net charge going out through this leg equal to zero. Otherwise, that can leave you exposed to corrosion through

this block.

28. Advanced Concepts for Charge Balance: Method of *Precision* via Feedback

So, there are two kind of research groups, research area approaches that are trying to fix this problem. Both use feedback. The first one, and this is an example from Sarpeshkar, is to use precision in a phase-to-phase method. So the idea here is to use techniques like correlated double sampling, auto-sampling of the current sources, to try to get excellent matching between the sources in the sinks in a pre-step, and then, when you're actually delivering the bi-phasic pulse, because of that auto-sampling, you have high precision and good matching. And that kind of, if you have good matching of your actual pulse widths and your dynamics, you can do a pretty good job of eliminating any net DC current flow into the electrode. That hasn't yet made it into any implant that I know of, but it is, it's a first order decent technique.

29. Advanced Concepts for Charge Balance: Feedback Loops for Polarization Management

The other one that's being explored is from Ortmanns and I'll point out that their group has a paper here on Tuesday, session 17, that I direct you to for kind of the "latest and greatest" in this area. They're using more of a long term, you can think of a servo technique, as opposed to demanding absolute precision within the sinks and sources. Over on this side, they're actually measuring what is the residual polarization in the electrode. Comparing that against a control reference voltage, and then, basically wrapping a balancing control circuit, kind of a global feedback circuit, around. That's shown here where you get your stimulation current. You have a residual polarization which you measure, and then you get some charge boluses to try to recover back.

And, so, there are a bunch of questions. As a circuit designer, you need to think about the system, how do we measure that ground, how do we deliver the boluses – a lot of details that need to still be worked out in a practical system. And that's, like you say, this is a paper from 2009 and they have a paper at this conference. These are the kind of issues that people are working out today. To keep you abreast of the latest developments, kind of the last thing to touch about, is there just a way to get rid of electrodes in a system?

30. Evolving Methods: Optogenetic Neuromodulation

And this is an area of kind of advanced research that our team and a bunch of others are looking at. And this is to go away from using galvanic stimulation to drive a neural network through an electrode, and, instead, activating the neurons with optsins and making them light sensitive and using that to drive the network. And, so there's a lot of research going on right now exploring the use of light to actuate tissue. It does add some interesting degrees of freedom.

31. Critical Interface: Optogenetic "Transconductors"

Beyond simply getting a stimulator that can excite tissue, what's really unique and makes us kind of excited from a circuit standpoint, is that there are optsins that actually inhibit neural activity. The inhibition of a neural circuit is quite challenging with an electrode, usually requires a lot of fast frequency stimulation or something equivalent to that. This is, this could, it's still very early in this field, but this could be a very disruptive technology in terms of its ability to interface with the body. The trick is that the power requirements are quite large.

32. Alternative Therapeutic Stimulation: Optogenetic Stimulator Prototype

When we built a prototype of the system and started to activate neurons and take a look at what's really

involved, so, you know, so, hypothesize, what does it take to build an optical stimulator? The current optsins and the light frequencies just don't, and the intensities required just in, and light scattering, and you stack it all up, it's really quite, quite challenging.

There's a lot of interesting technology work that can be done here, but a big part of it will actually be in coming up with better optsins. And, so the leaders in this group have now come up with some new optsins that are actually pushing into frequency bands that are starting to overlap with infrared, and so, that really starts to open up the degrees of freedom. Because now we're starting to get, you know you can think about it from a system design standpoint, all the gears could start clicking together. You've got an optsin that's sensitive at the right frequencies that then marries itself up quite nicely to an optical source that I can efficiently drive. So this is an area just to keep aware of, keep your eyes open for, because there's a lot of rich technology that can be brought to bear here as well.

33. Status of Actuation Methods in the Nervous System (2011)

So, I'm just kind of wrapping up our brief walk through actuation methods for the nervous system. This is kind of a snapshot in 2011 for just the nervous system for devices that are out there today. So ranging from the treatment of Parkinson's, to tremor, chronic pain, bladder, and that should be retention. I can explain to you the principles of operation of bladder retention after this talk. It's kind of interesting.

And then there's explorations for new areas. Clinical trials are going on right now for depression, epilepsy and fecal incontinence. So, the nervous system, so Steve Austerley who's, hopefully will get out of the blizzard in Minneapolis and make it tomorrow for the plenary, he has a saying that basically "Medicine is waking up again to the fact that we have such a strong electrical component to our bodies and to the nature of disease." And this doesn't even list some of the highest volume applications, which, frankly, are pacemakers, defibrillators and cochlear implants. Cochlear implants being, by far, the largest application in the neurospace.

34. Chapter 1: Summing Up Actuation Systems and Requirements

So, summing up actuation systems, so, it's all about management of energy, turning that into meaningful information for the nervous system and then driving it through the interface. If you want to be effective in this area, you need to, at least, get a working knowledge of the neuroscience for understanding the coding principles, and enough of material science that you can understand the polarization properties and constraints on your electrode system.

Everything that I've shown you, though, so far, is pretty much a static system. So we've just talked about stimulating the body, but how do we react to when something changes?

35. Presenting Problem: Actuation in Absence of Feedback …

And so this is the presenting problem for this chapter. The original pacemaker was pulled from a metronome circuit in Popular Electronics, a hobby magazine. Literally, Earl Bakken, the founder of Medtronic, was given a problem by a cardiologist: Build me a better pacemaker. Transistors were all the rage. They were kind of a new toy. Everyone was playing with them. And, so, Earl Bakken's genius was to say "You know, this system he wants me to build is like this metronome." So he built a two transistor battery powered pacemaker.

Here is, actually the first prototype, and, actually, the second prototype and the cover of the Saturday Evening Post with Dr. Lillehei, who was his partner. So the reason that I point out the second pacemaker is, I know it's the end of the day and so you want to have some entertainment, notice these little dials here, those

are actually sunk into the can, into the box. The original prototype had knobs, externalized. Of course, you put anything with a knob on a child, and what do they do, they turn the knobs. And, so, the first thing of systems design is how might people inappropriately use your technology. So they had to quickly get away from exposed knobs and sink them into the device.

36. Chapter 2 Sensor Systems and Requirements

But, those don't work so well for you. So, I've got a fixed-rate pacemaker. (sound of rhythmic clap, clap, clap, continuing through the next few sentences) So, it's good enough, serves you well enough when you're sitting in here listening to me, but suppose you're trying to fall asleep at night, your heart's going too fast, it's hard to get drowsy. Suppose you get up and the escalators break down for once and. gee, I actually have to take the stairs somewhere in the United States, well, now your heart's not going fast enough. So this is a real problem. And, actually, the second stage of our company's growth was to put an accelerometer (sound of faster rhythm: clap, clap, clap) into the device that measured your motion. And, based on your motion, it adapted the rate of the response.

So this kind of motivates the second chapter, I was starting to think about sensor systems and requirements for how can I actually measure signals from the body that might be useful for the treatment of disease?

37. What are the Signals for Detection?

So, the first thing that you ask, what are the signals for detection? Well, the signals can be both electrical or chemical in their nature. So if you, we have both kind of areas to think about. And, so, electrical signals, neural activity, muscle, drive signals like impedance, those are the first ones we want to start about.

I had this picture of the nervous system, and kind of with an Action potential shooting down in the channels, 20 nm, and kind of saying Scale Matters. You know there's a lot of action going on at different scales.

38. Circuit Scaling: Cells → Circuits → Networks

And, that's important because the scale at which you choose to interface with the nervous system really sets the requirements for what you're going to have to build as a circuit designer. You kind of, before you launch in, I don't want to see anymore ISSCC papers in our committee until I really say you've analyzed the scale that you're going to interface with the nervous system. Because you can go from microscopic to the mesoscopic scale, and believe it or not, neuroscientists actually do use this term, it's caught on in favor, somewhere in the middle of the networks. There's a 741 opamp. You can kind of think we had these little neural circuits that are kind of that level of complexity. And then, the macroscopic where you have the full systems on a chip. And kind of thinking about, so I, am I going to deal with the brain on a macroscopic level? So, you've got to first step back and say what's the scale of the nervous system that I want to talk to and try to derive a signal from?

39. Quick Primer: Origins of Bioelectrical Activity

And, so, I do apologize, this didn't translate. This is supposed to be a delta. I told Willie Sanson that this was a little box for you to fill in so you could be active. So they're supposed to be deltas in all these boxes that didn't transfer over on the slide. This is kind of the electrical model. Once again, kind of revisiting this membrane with inert potential. Most of the areas of interest for us are going to be dynamic signals. And so we're going to measure differentially the dynamic signal that's derived from the nervous system.

Remember when we talked about we kick off and we get a flow of sodium from the outside of the cell, or,

from the inside of the cell …… let me get this right. From the outside of the cell coming in, that's going to depolarize the Action potential and then we get a long-term rebound as the potassium moves and kicks back, and compensates the cell, and restores it. And, so you'll go home tonight, I'm going to spare you as you're trying to fall asleep and you're thinking about this tutorial and you're saying "Why doesn't it just wear down and everyone, like, drain your battery?" There are actually proton pumps that actually then, basically, actively charge and pump the ions back to restore equilibrium. So I thought of that once and it bothered me, so that's your answer. So, not shown on this slide is a proton pump that kind of goes and maintains this equilibrium.

So, differential signals, dynamic signals are what we're interested in.

40. Bioelectrical Measurement Principle (Bi-Polar Measurement)

If there's one thing I think this audience understands, it's how to build an differential amplifier. So we're going to go in and we're going to measure across two signals, amplify those and send them out. I grabbed this slide from Firat and kind of included it to add a note. One thing you have to think about is what's the reference bias in the system? What is the constraint on your common-mode? What might that do to your system? Many implantable signals are essentially floating. It can be much more challenging on an external system to get away with that. So, that's something that, as a circuit designer, you need to care about and say do I have a bias? What are my constraints in my application?

41. Cellular Scale: Origins of Bioelectrical Activity (What we observe with differential system)

So, now we get into the nature of the signals that we're interested in picking off, The origins of biological, bioelectrical activity. Starting at the smaller scale, so kind of in the region of an axon or cell. So, we've drawn a picture of an action potential going from the right to the left. So, the resting membrane potential then the depolarization, the action potential, and the repolarization. And this sets up current loops within the neural tissue.

And, so, if we look at a field potential externally, we kind of get the sombrero hat and if we take a first derivative of that with two electrodes, we end up getting something kind of looking like a bi-phasic blip. And, so, the time course of this, this is kind of looking at this from here to here, is on the order of a ms. So that drives the spec of kind of a kHz or so for picking off these localized high frequency blips, so a lot of the work from Reed Harrison and others kind of focused on this nature, the potential of where want to zero in and measure one cell's activity and that requires measuring these bi-phasic blips on the order of a kHz or so.

42. Ensemble Level: Origins of Bioelectrical Activity (What we observe with differential system)

For groups like my team where we're interested in more of a network problem. So we're not down at the granular level of a single cell. We're more interested in neurological diseases where it's kind of a network that's having a problem. Then we go up a level of scale and start to look at these broader loops. And so, when we're out, measuring at a more gross level of the field potentials or an electrocorticogram, an EEG, then we're starting to pick off a different signal. It's not necessarily the active potential's flowing. It's actually measuring the different flows of synaptic inputs and the return currents and getting these large dipoles. And, so, what we end up seeing is this large ensemble average of activity that kind of tells us what's the overall state of the circuit? What is the state of this room? So I kind of, you could have a conversation in this room, and, actually, right now it's not a very effective analogy so let's go to a stadium, so I might be turning to the person next to me, saying "So how's your 401k doing?" And they're saying "Well,the stock market's going great this year so I'm pretty happy." ….. has nothing to do with the game. So, that's not the level of granularity that we want to tap into. We're kind of pulling up one level, looking at this ensemble

average, listening to the roar of the crowd. And, that actually gives us some correlate to what's going on within the game itself. So, that's the kind of scale we're looking at and that's down at lower frequencies.

43. Bioelectrical Signals: Amplitude and Frequency Characteristics

And, so, Firat and, pulled together a nice summary slide from Webster from '92 and it kind of says this is my summary, what are the signals of interest, amplitude and frequency characteristics, and you kind of go from electrogram

Electromyogram, up these larger signals, and still fairly lower frequency. And I want to point out all the principles, even though we didn't talk about muscle, the principles are pretty much the same to the first order, in terms of the origin, it's just finer details.

Then we get into the nervous system and we get the local field potentials, ECoG, electroencephalogram. Those are kind of going in different scales within the neural tissue, but on the order of 1 to 10 to 100's of μV, but all at pretty much low frequency. And, then, the action potentials at that single unit cell, that's like what we were talking about, that's at the higher frequency, the KHz are above and somewhat larger signals in nature. So, these are the kind of signals that we're interested in.

44. Constraints: How Do We Sense These Potentials?

So, once again, you've got to think about, OK, I've got an electronic circuit, electrons, holes, carriers, but I, what these currents that are flowing within the tissue are ionic in nature, what constraint do I have in terms of transducing those? That's done at the electrode and that's some of the things we have to think about, the nature of our electrodes or the distortion that those can cause, this is not an ideal capacitor. The other thing is energy losses due to, I put polarization, because if you're stimulating a circuit, you need to think about the polarization recovery because that can actually be 100's to 1,000's of times larger than the actual signal that you're trying to pick off. So, managing the polarization is also part of building a good sensing amp. And, then, once again, safety considerations always come to bear.

45. Power considerations … What are the drivers?

We talked about power earlier, you know, what are the drivers? So, I'm going to start, actually, with the device longevity because we motivated that already. Our stake in the ground is therapy. Whatever it takes to derive therapy kind of sets the power budget for the overall system. For sensing circuits, we like to have things on the order of equivalent, or more preferably, down at kind of 10%, it depends on how you're going to use the sensing.

If it's to build a closed-loop system, you want to have it down, of course, less than the therapy because you don't want to actually, unless closed-loop is required for effective therapy, you don't want to actually undermine the device longevity by adding a sensor. For diagnostics where you're only providing additional information, the energy budget is actually pretty small because you may not be materially improving the outcome of the therapy.

The other one that you have to keep in mind is the biophysics of the tissue. You can't just drop a watt in your chest. I mean, you've got, kind of that intuition. We do have specifications from the regulatory bodies that you need to keep any temperature rise below a few degrees. That's kind of an absolute max. You've got to be under a few degrees centigrade increase. Reed Harrison, well, Utah did a nice little study for his implant array and kind of derived 10 mW for the implants that are going up over your motor cortex. So, that's one example.

So, a lot of the amplifiers that you see at ISSCC and stuff, when they give the power requirement that goes back to that analysis for motor prosthesis. They're not limited in terms of their longevity because they're all using inductive telemetry systems, or inductive power systems, that are always there. So, you've got to be careful because you say that "Oh, I can use that no matter, whatever power that I need, it's available." You've got to think what are the other system constraints that come to bear. In that case, it's to avoid any tissue damage due to heat.

46. Application & Constraints Define Specifications

So, here's kind of, now we get a little bit closer to integrated circuits. Kind of the application constraints and specifications for these different scales. We talk about EEGs, ECG, ECoG, local field potentials, Action potentials. And, one of the challenges that we actually find in this space, in this design space, is that it's too new. So, there actually, if I design an ECG amplifier, I know when I'm done. There's a set of international specifications. You meet those specifications and you can move forward with the regulatory body.

The area with ECoG, field potentials, the arrays that are used, there just not that well developed. It's still very much still in the research stage. It's an opportunity for us, but it can also be a challenge to know when you're actually done in your application. We're going to re-visit that at the end of this talk.

These are all quite variable. The one thing that was missing in your sheet that I threw in when I was doing my practice was realizing that I didn't specify a bandwidth for the noise, and, so within these lower frequency signals, ECoG, LP and such, these are kind of in the 0.5 to 100 Hz range roughly. While for measuring action potentials, you need to keep on the same order of μV_{RMS} while boosting your bandwidth up by up to 2 orders of magnitude from what you have in an ECoG. So, that's why these systems tend to pull significantly more power than an ECoG system, nothing magic about it.

47. Review: Industry Standard Instrumentation Amplifier

I'm going to give you a very brief tutorial of kind of the instrumentation amplifiers and thoughts that are brought to bear in this area of measurement. This is just to give you kind of a background of what's raging. In fact, I'm very happy the plenary tomorrow from IMEC makes it really clear we don't have the golden instrument amplifier yet. So, they're kind of, which surprised me, because in fairness, you know, I actually respect them, I thought they were going to show IMEC as the gold standard. But, they've kind of showed all the constraints and degrees of freedom. This is still actually a pretty interesting area of exploration. So, this is just a snapshot of the "the state of the art".

The industry instrumentation, industry standard or A.K.A. this is what we learned in school, you know, kind of classical three transistor, or three operational amplifier, instrumentation amp, lots of power in the amplifiers, lots of power to drive the resistors not commonly used in this space because of our constraint for low power.

48. Instrumentation Amplifier Overview (Resistive Gain Element)

Toumazou came up with a nicer version of this with the instrumentation amplifier presented back in '89. The idea being that we can have a buffer, drive it across a resistor, then mirror that, the current flowing through that resistor over to a secondary resistor to give us the gain, R_2/R_1, and buffer the differential output. And, so, kind of the input impedance characteristics are limited by the parasitics that we see and then you just have to be a little bit mindful of the excess noise that you might get because you have two buffers and the resistor before you have any gain. We're actually going to re-visit this with the example from IMEC and their instrumentation amplifier.

49. Instrumentation Amplifiers Overview (Resistive Gain Element)

You can get pretty good power dissipation, and you just have to be mindful of the limitations of your buffer for a common-mode rejection ratio. The other thing, you know we talked about that polarization headroom. That can give you DC potential across the electrode that you need to manage. So that's a problem with these kind of direct coupled amplifiers is that you might need to have some kind of blocking circuitry to deal with the polarization that you see. Because those sensing electrodes, like we were talking about, those are also acting as polarizable, basically batteries, that can superimpose onto your signal.

50. Instrumentation Amplifier Overview (Capacitive Gain Element)

The other extreme, extreme is not right, the other kind of approach is to use capacitors as the gain element instead of resistors. I think it's fair to say, Harrison's 2003 is kind of a de-facto example. It keeps popping up. I think part of that reason is it's almost a kind of a textbook kind of design where we have coupling capacitors coming in, a feedback capacitors on an op-amp to give us our gain. And then, of course, we can't just have capacitors, we need to have a resistor in this system for its bias stability, and that was achieved through sub-threshold transistors.

So, pretty good gain, relatively large input capacitance, or impedance as long as you make your capacitor small. It's usually on the order of 10 pF or so. And, a pretty good noise characteristics.

51. Instrumentation Amplifier Overview (Capacitive Gain Element)

Kind of the only, the trade-off, a bit, is the amount of CMRR you can get practically in a typical process for the size of the capacitors used. But, you do get good power dissipation and you get outstanding DC headroom because you're actually coupling, you've got an AC coupler right at the front-end of this circuit. By putting those sub-threshold transistors in the resistor, you can still keep that AC cut-off point down in the order of Hz or so without too much effort.

52. Operational Transconductance Amplifier

So, in, kind of getting a little bit into the design of that transconductor, just a few things to point out that are worth noting when you're studying these designs. The trade-off between bias points for the different transistors is read, you know, on this input transistors, put them into sub-threshold weak inversion to maximize the g_m/I_D ratio and then all the other, so, this is a pretty simple transconductance op-amp. So, all the other current mirrors that are used, these would be pretty hefty noise sources, except he's run them in strong inversion and so that the relative transconductance referred back to the input is somewhat small, so that keeps the noise in check. Now, you also, this design is susceptible to 1/f so you've got to size these input transistors quite a bit. And, of course, that hits a limit with these input capacitors and feedback capacitors. So, there is sort of an optimal design point that the designer can identify.

53. Measured Neural Amplifier, Input-Referred Noise

So, works like a charm. It's a good amp and I think that's why versions of it keep showing up everywhere. Especially, it does a great job in the region of the action potentials where you want to kind of measure on the order of KHz or so. You do see at lower frequencies, kind of where we're interested in EEG, ECG and ECoG is that you, pretty much, are you're dominated by 1/f noise. And, so you're more, you can be limited by the transistors, your process and things like that.

I don't want to over-play that because if you actually look at the spectral density of the brain, it has a 1/f

noise profile. There have been some kind of nice papers where they actually bias the 1/f noise of their amplifiers to the 1/f noise of the typical brain and play tricks with transconductance. And, so, it's, don't, let's not do over-kill on the impact of 1/f. I think you need to be thoughtful as a designer because the system that you're hooking up to has a pretty significant 1/f noise floor. There's actually a lot of interesting biophysics about the spectral noise property of the brain that's worth reading about. But, the reason we've explored it at Medtronic, in terms of looking at things like chopper stabilization studies is actually more concerned about process variability and not as much about just the inherent 1/f noise. We talked about that in an old ISSCC paper.

54. Dynamic Compensation: Chopper Modulation, Address Several Amplifier Non-Idealities

So, if you do want to go and try to become a little more process immune and deal with the 1/f noise floor, chopper modulation is a nice way to go forward. We have Kofi Makinwa here so, what can I say, I should give you the floor for ten minutes. The principle of operation is we take our input signal and then we have an amplifier with quite a significant offset in 1/f noise. And, then, what we do is, before we present the signal to that corrupted amplifier, we're going to up-modulate it with a chopped frequency, with a modulator which is just, basically, in most instances, a simple across bridge of CMOS switches.

So, the signal gets up-modulated, then the amplifier amplifies it, also corrupts the signal chain with the offset 1/f noise, but, then, on the outside of the amplifier, we demodulate the signal back down, and the 1/f noise gets sent up to the carrier, and ideally, our signal comes back down to baseband and we can lowpass filter out the 1/f noise. So, that's the general principle.

55. Chopper Modulation, Compensation for Amplifier Non-Idealities

Some of the advantages is, of course, the 1/f noise is dealt with. There are also a bunch of mismatch related errors that can also get sent out of band. This can help our CMMR a bit. The disadvantages, you know, we do kind of undermine our input impedance a bit. You have to be mindful. We now have capacitances switching at fairly high carriers on the order of KHz and you've got to worry about what does that do in terms of the input impedance of the signal. As drawn, like we were talking about before, any polarization on the electrode is going to get sent into this signal chain and might actually saturate it. So, you need to deal with that as well.

56. Chopper Modulated Instrumentation Amplifiers (Resistive Gain Element)

So the, kind of within the resistive gain element, kind of a nice way to go about just chopper modulating is to put the switches onto the input and switches onto the output. And then, use the internal operation of the instrumentation amplifier as is. The one concern you just have to keep in mind is the need to consider is what is the bandwidth of the system in the micropower domain because if we're chopping too fast for this system to respond, we can actually end up getting some parasitic dynamic performance that needs to be considered. That can be dealt with and we'll talk about that here in a second.

57. Chopper Modulated Instrumentation Amplifiers, DC Blocking using a Servo Loop

The other issue, in terms of the polarization voltage, can be dealt with a little servo loop. And, so, both Firat's team and then our team kind of explored variations on this. But the idea is instead of having that input signal come in and just be amplified with the electrode offset, as we send it through the system, we get all the benefits of getting the 1/f noise, but then what we do is wrap this DC servo around. So, we measure what's going on at DC, kick that back over, kind of extract that signal, and ideally with an integrator, we can then up-modulate that and superimpose that back into the amplifier and strip off the offset, so as shown here.

That doesn't come for free. We've also stripped off the lower frequency information from our signal. But, in a lot of applications, the polarization signal is changing at a much lower frequency than the signals of interest for us. But, we have to be careful because that's actually the classical way that you'll never discover that there's signals of interest at those low frequencies. So, always be careful about going falling too much into religion and not keeping an open mind on these things.

58. Chopper Modulated Instrumentation Amplifiers, DC Blocking using a Servo Loop

So, this is the, IMEC came up with an amplifier so you can see the, let me point out the core operation is the Toumazou amplifier where we have, think of these originally as just current sources. We have a servo that basically measures the current that's drawn across these, sorry, let's just step back, the PFETs act as the buffer. We then drive current through this resistor, R_1, based on the difference between the two potentials, through this servo loop, then we impinge that current and drive it across R_2 and get voltage gain, R_2/R_1. So, the way to chopper stabilize this is to modulate the inputs and to modulate the outputs.

You know, I talked a little bit about the dynamics of the system. What's nice about this is there's one dominant pole pair, right here, in terms of a high impedance node. So, they can still get pretty good performance even at low power. And the other thing that I actually really like is that because of the nature of this PFET and that they drive and servo to keep a fixed current through it at the chopper frequency, it has incredibly high input impedance. So, they've actually touted this and it's very impressive in terms of the input impedance of this amplifier. And, so, in terms of dealing with the polarization, like we were talking about, they then strap a feedback loop around. And in terms of their high level biasing, kind of add a little bit of wiggle on top to compensate for the DC potential that's coming in. This has been quite a successful amplifier. I actually really do like it quite a bit.

59. Chopper Modulated Instrumentation Amplifiers (Resistive: Baseband vs. External Modulation)

And, so, the, kind of comparing it against its baseband, you know, they do get the benefit of much higher CMRR due to the chopping, much better noise performance, we're getting rid of the 1/f noise, modest input impedance degradation because of those benefits of having a buffer architecture like we talked about, very modest power penalty, and then, there is just a finite DC headroom that they can account for. But, this is a good design. It's quite powerful.

60. Chopper Modulated Instrumentation Amplifiers (Capacitive Gain Element)

Going toward the capacitive gain element and trying to patch that up. There's some architectural choices that can be made. You know, where do you chop?

One approach is to go into the OTA and just chopper stabilize that. And, then, try to just basically clean up the OTA because the capacitors are the first order ideal elements.

The other is to go outside of the capacitors and to modulate there. Of course, to design so there's trade-offs with each.

61. Chopper Modulated Instrumentation Amplifiers (Capacitive Gain Element, Internal Modulator)

So chopping within the amplifier, this is the Verma circuit, and, so tomorrow at the IMEC plenary, this will say MIT, this is when Naveen was a grad student at MIT. The idea here is that they chop at the input and then within the amplifier, kind of basically taking the folded cascode paper that we did in 2008. So, that's just kind of moving that over here and putting the chop switches there. The trick is that, you know we were

talking about you've got dynamic switching in a parasitic capacitor here, as well as the gate capacitance, you know these aren't acting as buffers. So, any DC voltage across here, this is going to look like a finite impedance. And, so, to kind of keep control over that, they are using off-chip capacitors.

So, these are the way to keep your reasonable high pass is to have a huge capacitor, and these are off-chip and so you get some pretty modest matching. It's tough to get 60 dB, at least from our vendors, and 60 dB doesn't really cut it in a lot of applications. So, I'd say the areas for improvement, for those of you interested in this area, is how do you get the performance benefits of this system without the need for off-chip capacitors

62. Chopper Modulated Instrumentation Amplifiers (Capacitive: Baseband vs. Internal Modulation)

which is what I kind of drew on here and that's why I circled it in red. And then the other thing with the input impedance, we've got to be careful because we'll say it's high, but if your input capacitors are off-chip and large, the actual effective input capacitance can be quite small. So, that means our input impedance will also be somewhat small. So, room for improvement in all these amplifiers, those are two things that I think about when I look at this.

63. Chopper Modulated Instrumentation Amplifiers (Capacitive Gain Element, External Modulator)

The amplifier that we have, which people say Medtronic, I wouldn't say that, we don't publish everything, of course, that we do, but the one that we've publicly talked about is the one shown here. And the idea here is we kind of chop on the outside of the capacitor and there are some, the biggest hit I think here, is taking a whack at input impedance. It does cut that down to the order of 10s of MΩ or so. And then we've implemented our DC servo loop with the same principles we've already talked about with this shunt feedback capacitor through a smaller capacitor and an integrator in the servo loop.

So, it does clean up the 1/f noise. Then you just have to decide how much headroom you need to feedback because, of course, as you make that capacitor larger, it adds more capacitive load to this network which can come at a trade-off to your noise performance. You've got to kind of design the whole system based on your constraints.

64. Chopper Modulated Instrumentation Amplifiers (Capacitive: Baseband vs. External Modulation)

So, the, like you're saying, the biggest thing that makes me sad on this one is that we have good CMRR, it's never infinite in an ideal world, but, you know, we can get up to the 110-120 dB with our system. The input impedance is where you take the biggest whack. The thing is, we do kind of have a decent estimate for what this chop effect is on these capacitors and we know from a differential standpoint, we have a measurement of that, so one can do an estimate and try to do a little positive feedback loop.

And then your reference guide, Kofi Makinwa's team took a whack at improving the input impedance with that positive feedback network. They did a nice job of also giving its limitations. So, there's still work there. I've historically said you could do the servo loop and then, if you really cared to, and then Kofi said "yeah, well, but…", and so, if you want to know the "yeah, but", I refer you to this paper.

65. What are the Signals for Detection?

So, the other thing is, kind of stepping back and I think Yoo tomorrow will give you a nice feel, it's, there's a lot of interesting area, the amplifiers have not been worked out. There's no gold standard right now. The device, the amplifiers that we're looking to put into devices actually have gone and compiled some of the best

of the three amplifier architectures I've shown here. There's still a lot of interesting work that can be done. So, I'd encourage you to go to session 17 on Tuesday. There's a couple of amplifiers being presented, including a nice one, in 65nm with some interesting tricks. So, it's definitely still a rich area.

The other thing that's kind of coming out is this idea of chemical sensing. So, we've always been biased in thinking that, at least, I've been biased and my friends at IEEE for measuring electrical properties, but, we have to remember that chemical transmitters are actually a key computational property of the nervous system. One might be able to measure those transmitters and learn a lot about the network. And, so, how might that be done?

66. Making the Electrode Work Harder for You?

And, so, this is highlighting some of the work by Pedram Mohseni's team, it's been presented here at ISSCC, printed in JSSC. From a high level point of view, you can think of it philosophically as making the electrode work harder for you. We talked about electrodes that have chemical reactions, well, what if we actually used that electrode and sensed specific chemical reactions? That can tell us something about the nature of the fluid surrounding the electrode.

67. Emerging Application: Fast-Scan Cyclic Voltammetry

And, so, what their team is doing is going in and designing systems that sweep the voltage, so, kind of, you can see this is the fast-scan, as the relative frequency, up, down, cyclic voltammetry, so going in cycles. And, what you do is basically you drive a voltage across electrode and measure the resulting current. And, so, I give them credit because it, I certainly, the first time I would look at this graph, I might say nothing happened.

But, what they do is measure the background total current, do it against a current reference, and then they can actually subtract the two and look for a resulting faradaic current. That resulting difference is basically giving them information about the electrical reaction at specific voltages. And this can give them a signature, so they can actually look and try to pick off different chemical signatures at different voltages that are driven and use that to tell them a little bit about what's the chemical environment of the neural tissue that's surrounding their electrode. It's quite an interesting degree of freedom. It tells you more about just the electrical activity, but also the chemical environment that the neural tissue has around it.

68. Time-Share Voltammetry and EPHYS

And the reason I pulled Pedram's work out in particular is that he's published on some systems that say instead of doing one or the other, let me time-interleave the two systems. And, so we can go in and have a phase where we're measuring the chemical environment of the neural tissue, and then we can transfer over and inter-phase where we're going to listen, sort of passively, to the activity within the nervous system as well. He's looked at building systems that can do both and reported on those results.

This is kind of an interesting area of future endeavors which we watch quite carefully. There's also a team at the Mayo Clinic that's been doing quite a bit of measurements with this trying to build a closed-loop neural stimulator and trying to work out the principles of that using chemical measurements as well.

69. What are the Signals for Detection?

So, the last kind of area for the sensing to talk about are these indirect measurements of patient state. And the idea here is, kind of give this some context, the first pacemaker that they wanted to make adaptive, they

brought in all kinds of cardiologists from around the world, put them into a working group for several days and said "What are the signals in the body that we should interpret and then use that to drive a pacemaker and change its pacing rate?" Many pages were typed up, things ranging from we're going to pick off nerve traffic on the vagal nerve and use that, interpret it, decode it, and then drive the pacemaker.

The one that actually got a lot of interest was listening to the heart valves …. glic, glic, glic, glic ….. and heart sounds, and using the heart sound coming from a device to drive the pacemaker. So, they actually were building a prototype of a heart sound device and kind of measuring the valves and the sounds of the heart. And then the prototype, in the notebook, it's actually quite famous now at Medtronic, the scientists said "This might work if the damned dog would stop moving." The whole point was they fit this whole system to measure acoustics, but, actually, it was a great motion sensor as well. And then the scientist, who's now quite high at Medtronics, said "Hmmmm. You know what, motion is kind of what I need my heart to increase. Let's do a very simple servo loop."

And you can imagine actually, it's not very pure, it's kind of an engineering hack approach. There's a lot of resistance to it, that's not physiological, that's not representative, but it was actually a pretty simple way to build a good loop. And so, that's why I'm kind of teasing out these other "indirect measurements" of patient state because engineers have a lot of tools that we can kind of make indirect correlations and build meaningful systems.

70. Sensor Fusion with Accel, Surrogate Markers for "Brain State"?

So, one of the areas a lot of groups are looking at is strapping accelerometers onto people and doing things with them. Accelerometers are kind of ideal for an implant. They build off a lot of predicate work with pacemakers, rate-responsive pacing and you get into some philosophical questions that are quite difficult of how does one measure pain? What are the measurements of depression? We don't have very good correlates right now, like, within neural Activity.

So in the absence of that, people are looking at things such as sleep patterns. So across the X axis is time and days, across the Y axis is an activity count. In a control population, you can actually see this nice circadian rhythm, eight hours of fairly low activity as they sleep through the night. But, in a depression sub cohort, they're actually showing poor sleep quality. People are trying to look at these indirect measurements and saying "Can't I build a reasonable diagnostic?"

71. Implantable Accelerometers Possible?

And, so the other area I would suggest you kind of track is looking at implanting accelerometers. The reason for that is the technology has come to bear in the last decade, thanks largely to consumer devices and automotive accelerometers, where we can actually build pretty reasonable accelerometers for implantation. Things that are operating on the order of μW and can pick out signals that are of physiological interest.

Out of my own curiosity, kind of thinking about figures-of-merit, you know, the problem with accelerometers is it's all dependent on what sensor you hook up to it. And, so if you want to look at the design, you need to, somehow, kind of, if you will, normalize out the interface circuit and that's what I'm trying to do here with this simple figure-of-merit that we discuss in this paper. You kind of need to do that.

The last two papers are from the recent ISSCC, especially this last paper has a very impressive noise floor. But, when you work through the figure-of-merit, you kind of realize, well, a lot of that low power is because they have a very compliant high capacitive per g front-end capacity, you know what I'm saying, front-end capacitor, so they get a lot of signal coming from their MEMS element. And, so, kind of keep that in mind when you're looking at different designs. But, this is another area to keep in mind of, you don't necessarily

have to have a direct measurement of an electrical signal or a chemical signal, but you can also use something that's more abstracted out.

72. Chapter 2: Summing Up Sensors and Interfaces

So, in summing up this kind of area of the sensors, these are kind of the bigger areas of actuation and sensors. Signals are low bandwidth, they're prone to confounding inputs and things like 1/f noise. And, probably realistically in the future, we'll need to "fuse" together multiple sources to get any kind of specificity in identifying patient states.

So, detecting signals is easy. How do we manage the data?

73. Typical Presenting Problem: Neural Signals Require Amplification and Transmission

So, here's a presenting problem from my, Reed, my friend, Reed Harrison. He builds these arrays; he puts the amplifiers on top; he's measuring action potentials, these are coming forward like we were talking about, order of ms.

74. Multi-channel neural recording devices produce lots of data!

This information comes off all these electrodes, 100, suppose you wanted to actually digitize everything and send them across the boundary, you end up with 15 Mbit/second telemetry rates. So, just raw data streaming, this is a big issue.

So, the next method, which is what Reed was exploring, is to put a simple threshold comparator in, and then what you send across is basically just an event packet. So, some send out, as opposed to 15 k, go down to 1kSamples/second resolution on event packets, 1 bit spikes instead of 10 bit signals of data and you start to drop down to more reasonable information rates for a micropower implant.

So, the point is not to say, this is the solution, the point is to say, you've got to think about how you're going to manage this information. And, that's the point of this section.

What are algorithms and data extraction?

75. Traditional Biopotential Signal Analysis Flow

And, what are people working on? So, I kind of took one of Firat's figures from the paths where we had data extraction. Where we take our signals in, this is an example of a heart rate, we're going to amplify the signal. Then we get into feature extraction where we're going to say, OK, there's something going on here. We detect it as an R wave and then we classify it. In this case, it's like an atrial fibrillation. The clinical impact increases the higher we go up this chain. So, what we want to do, and that I think one of the challenges of many of our JSSC and ISSCC papers is we kind of stop here and that's why we don't get as much of an audience because, OK, you built an amplifier, so what. You know, how do we actually finish this chain and push it farther up the value curve for eventual customers?

76. Tension: "Optimum" Point for Digitization?

This gets into some philosophical debates, some that we'll talk about through this conference. There's what is the "optimum" point for digitization? So, I have to, I disclaim the choice of words, this is Rahul's slide. He likes to say it's meaningless high bandwidth numbers at the ADC. What I'll just say is, a lot of bias is

digitized very early in the signal chain and then throw it to a DSP. And, when we do that, this can actually be quite power-hungry.

The other extreme is, say we're going to go and do as much analog pre-processing as possible, do the extraction, do the ADC, and then go to a low power digital analysis. How do you partition your chain is a big part of the system problem.

77. Principles for Ultra-Low-Power Design

So, one of the nice analyses I think that he did do in his 1998 paper was to say "What are the relative trade-offs between digital and analog processing, both in power and in area?" And, kind of, you see an inflection point. So, this is, I put this down as kind of a rule of thumb, not an absolute, but below 60 dB signals there actually can be an advantage in the power and the area trade-offs analog pre-processing for some simple functions. And, so, it, though, a designer really needs to take these into mind and kind of say do I need to digitize this early in the signal chain, or can I do some analog processing before that.

78. Required SNR for Biopotential Applications

And the reason for that is if we go back to our biopotential applications and kind of think it through, the signal to noise ratios for many of our signals of interest, like ECG, action potentials, EEG, they're actually all down at the level where analog and digital can kind of go head-to-head with each other.

I did add this, thinking about it, when you start stimulating the presence of these signals, you can get up to 120 dB and then digital is actually out of the question, too. You actually need to step back and, but, that's a tough problem. And, that's for the follow-up tutorial next year, how to deal with sensitive interactions.

79. Consider a Typical Use-Case: Cardiac Filtering

So, here's a typical use-case to kind of tease this out, cardiac filtering as an example. And, so we're going to, we've got our signals, so here's amplitude versus frequency. These are the signals of interest, the P wave and the R wave. So the P wave is that first little blip on the complex, and then you get your Q_{RS}, and then your T wave is that kind of rebound effect. What we generally want to do in an algorithm is pick off the P and the R waves. We want to reject the T waves because those can confound a detector, also know, look like your heart's beating twice as fast as it is. And, we also want to reject muscle signals. So, within this constraint of the cardiac realm, we want to pull out this frequency of information, extract it and identify the signal while rejecting these out of band signals. And, very similar requirements and thoughts go into brain sensing and EMG.

80. Analog Signal Processing, "Power Efficient" Frequency Analysis

There are multiple ways to go about this. Our team is taking a look at trying to do it in the analog domain, kind of merging those concepts of a chopper, amplifier, and using the kind of natural frequency multiplication that's inherent within a chopper chain to extract the spectral density. And so, we can get basically with two chains an in-phasing quadrature signal and sum them together with a low pass filter and get a relatively efficient frequency analysis. I will say this is an approximation and we documented that in our, this reference here is just how much of an approximation is it.

81. Chopper + Heterodyning Approach

So, that's the mathematics, kind of the intuitive view is by shifting the clocks within the chopper amplifier we

can pull out a specific bandwidth of interest, in terms of the center band through this delta modulation, and then we can pick out the effective bandwidth around that center frequency with the output low pass filter. And then we tend to add one additional low pass filter to derive the net spectral envelope coming out. So, very simple, very simply like an AM radio hooked up to the brain.

82. Improved Analog Frequency Analysis, Example: Application to Cardiac Signals

So the IMEC group improved upon this in 2010 by basically realizing that the input stage could be shared between both the in-phase and the quadrature signals. The other improvement they made to the system was to digitize the in-phase and quadrature, and perform the summation of the power signal in the digital domain. And if you want to know why they did that, when we did ours in the analog domain, in order to get μV kind of signal resolution, you've got to do nested-chopping. You've got to chop everything. I mean, there's choppers everywhere to try to get this thing to pick off μV. To Firat's credit, they said just digitize it. It's back into that slide of the best point for digitization probably is back at this point in the signal chain, don't try to be a hero in the analog.

83. Power Efficient Frequency Analysis, Example: Application to Cardiac Signals

So, they presented their work in JSCC last year, built a very nice system on a chip and built basically a cardiac detector and implemented it. But, now we do the counter-point.

84. Now Consider the Counter-Point!

So, of course, we're never doing anything static. Kwong and the MIT group at ECirc just this year presented their digital processor. Starting to put in some very interesting sub-blocks in terms of FIRs, FFTs, the square roots, some non-linear things that, just the sorting data, the median filters. And so, what's happening is that digital signal processing is giving analog a run for its money.

85. Signal Processing Platform Overview

The reason I say that in terms of my space is that remember we do have a stake in the ground. So maybe digital won't always be the lowest, but you also have to put on your systems engineering hat and say is there a point in my system where I get a "don't care"? So, if you tell me you can do a bandpass filter for 5 nW, but I'm building a stimulator that has a background stimulation level of 100 μW, you're not really saving me that much work. You know, so, then you've got to look at what are the other benefits of digital signal processing in terms of flexibility, in terms of design for test and other areas. So, I think this is a very rich area for debate, and it will actually should be a lot of fun over the next five to ten years.

And, so, kind of the approach the MIT group took was to pick out a node that had reasonable leakage in power, and then to build dedicated blocks. And so they had the best of both worlds, in that there is kind of a standard micro-core that they could run with, but then they could pass it to specific sub-units of interest for biomedical signal processing, particular FFTs, CORDICs and median filters.

And so they worked out two examples

86. Application 1: EEG Feature Extraction [1]

to kind of show the performance. One is an EEG feature extractor which is for seizure detection. Ali Shoeb was one of their co-authors and actually wrote a nice seizure detection paper. And then Ali actually did a little stint with us and talked about some of the constraints.

The high level ideas, they want a bank of filters to go and extract different band powers and then they pass this off to support vector machine. What the MIT group was showing is that by going into a dedicated accelerator for extracting the band powers and looking at FFTs, they could actually save about 10X the power of their detector. I have to say for the amount of complexity in their system, this is actually a pretty compelling number to get out of it. The one thing that's missing is that you do need to add on your feature extraction output. You've got to show what is the energy support factor of your machine.

87. Application 2: EKG Feature Extraction [2]

The other application is the ECG, electrocardiogram, and in this case they did the low pass filter and then actually looked at the line length going through and then combining that with adaptive threshold, could say do I have a QRS complex or not? And so, they did see a significant performance improvement. What I would be interested in is, this is actually still quite a bit of current for an ECG monitor, that's still quite a bit spindy. And, so, you get into the last kind of close out of this discussion is, how do we ever figure out which one is right? We get all these circuit designs coming in and they all have different specs; when are we done?

88. Having a Run-Off? Validating the Signal Chain

So, having a run-off, so this is actually something I'm going to challenge us as within the IEEE community for this conference is the validation of the signal chain. So, regulatory requirements don't say that verification is enough. What I mean by that is there's some nuance language. Verification is saying did we build what we say we would build? Which is kind of easy to execute. So, that's things like what's my noise floor? What's my bandwidth? What's my power?

Regulatory bodies require you to validate it. Which is did you actually build what **needs** to be built? And that's one level up in the abstraction and so it gets into things like are you actually able to detect heartbeats with your system? Are you able to detect seizures with your system?

And so, I actually worked with Ali Shoeb and put these in your book. These are a whole set of references that are available on the web for free that have been annotated by clinicians. So, if you want to actually understand the performance of your system and compare it, you can actually download these different signals and run them through the ringer. And, we did add on the bottom TEST on ALL the DATA because a lot of times folks will actually train and test on snippets of data and say that this is absolutely fantastic. Trust me, run it on all the data. You will fool yourself otherwise. So, we meant to have this as a reference. So, I'd encourage you, go off and validate your signal chain.

89. Example: Preclinical Validation of "Brain State" Detection

So, here's an example to show, kind of, what do we mean by validation? This is taking our brain processor chip from a couple of years ago at ISSCC, which is now in a research device. And, so, what we want to validate is basically, the capabilities of this processor equivalent to what can be done with off-the-shelf instrumentation. And, so in collaboration with Dan Moran down at Washington University , what we actually do is externalize electrode systems from a non-human primate that's controlling a cursor by thinking about moving the cursor. And, then we put our system in and extract a signal. And, then compare it to off-the-shelf equipment.

90. Typical Sample Run Real-Time Cursor Control

And, so what we do, this kind of gives you a better feel, is the

primate's given a task. It's shown the screen.

It's thinking basically modulating its brain pattern activity. Our system detects the ability to go and put the orange cursor over the green arc, kind of a binary task, looking at the accuracy and the time. And then we go to an off-the-shelf system. Basically, you go back and forth between the two and use that to derive

91. *Quantifying* Detection Performance

the receiver operating characteristics for the system. So you say "what was the signal that was going left, what was the signal going right?"

We can then build our algorithm and simple detectors between those two states for driving the cursor,

looking at the trade-off points. And, then actually run, head-to-head, our system running in its low power mode versus a control system running off basically 100W plugged into the wall. And this gives us something that's actually arguable. It's tangible. So, going into these validation steps, actually running things head-to-head, it tells you how much, how well do they really do in a representative application. So, I'd encourage everyone to start, kind of as a next level of maturity within the ISSCC, is let's start thinking about doing these head-to-head comparisons and representative applications.

92. Chapter 3: Summing Up Algorithms and Classifiers

All right, so summing up. I'll give you one example and let you go. So, algorithms and classifiers, analog vs. digital signal processing vs. telemetry. It's a very rich field. We're actually at a great time because they're so close to each other that you have an incredible degrees of freedom as a system designer. But, in order to actually make some forward progress on this, we've got to get into this validation mindset. We need to kind of come up with a methodology of validating that our solutions work and coming up with a more quantitative way to put designs against each other in the future.

93. Chapter 4: Putting it Together! State of the Art Systems Examples

So let's talk about one example of how we put these things together in practice.

94. Status of Actuation Methods in the Nervous System (2011)

In a representative example, from chronic pain, and so this is within our area.

95. Spinal Chord (sic) Stimulation: Unmet Need

So this is the unmet need, where's the video? Oops… sorry), so in a stimulator today, the electrodes are actually fixed with respect to the spinal column, but the spinal cord is tethered on the side. So as patients change their posture, the volume of tissue activation is actually dynamic. You can kind of think I've got a fixed electrode, fixed stimulation, as my spinal cord moves, I'm going to couple different amounts of energy into the column.

And, so that's a problem because when we look at the net system

96. New System Architecture

the patient basically needs to be given this programmer. And, so that's what's shown here. So they're

constantly adjusting their stimulation, so I have it here because it's a buttock implant. And so, if they're going to lay down for the night, they change their stimulation. If they want to roll on their side, they might need to change their stimulation. If they change their activity, they might need to change their stimulation. So they are actually the sensor and the algorithm for their device. They are closing the loop.

So using that accelerometer that we discussed earlier, we get a nice feedforward mechanism. We can actually measure, detect the posture, detect the activity, throw that into an algorithm, and then, we can, through the implant embedded within it, go through and close the loop with the goal of taking that burden

of the programmer away from the patient

97. Technical Solution for Unmet Need

So this is kind of that same schematic picture. Now as the patient is changing their posture, making adjustments, the accelerometer is detecting that change in posture, sending a signal to the stimulation engine, and then it dynamically changes its amplitude appropriately. So, believe it or not, very simple system. It's currently CE marked and under investigational use in the United States. And so it's kind of one example, pulling together these principles that we've talked about of actuating the spinal cord, but then adding sensors and algorithms on top of it in order to build a very nice system for the patients.

98. Our Other Mandate: Enabling the Next Generation of Scientific Discovery

And then, the last thing, which I want to just put as, I think, our other mandate as engineers, is enabling the next generation of discovery. So a lot of the papers that come through ISSCC in this field are about scientific discovery and this is important. So the schematic I have up here is a sheep. The AN represents the interior nucleus of the thalamus, the HC represents the hippocampus. And what these are tapping off of, what these two electrodes, the 3387, the 3389 is a circuit of papes which has been implicated in some forms of epilepsy.

And the idea, the reason that I put this up here is that that research device that we were talking about with the cursor controller, has now been implanted into some sheep models. And what we're doing as electrical engineers working with biologists, is understanding the dynamics of that circuit so we can get a better feel for how is the circuit working? How is it not working? How might we as engineers build a control system that allows one to modulate the network for the treatment of epilepsy? And so, it's not just about building necessarily a product today, the other thing that the neuroscience really need for the future are the building block tools, that they can do the next generation of research with. And so I challenge you all to help us develop those as well.

99. Closing Summary Thoughts …

So, in closing, it's all about the system as far as I'm concerned. The best designs take a holistic systems approach. The actuation, I think there'll always be a strong bias for actuation systems because they can oftentimes address issues with technology. It's also one of the areas that's really developing quite quickly. And we talked a little about optogenetics and the need to remove capacitors. There's the paper in section 17, talks about new "stim-engines". There's still a lot of opportunity there. Sensors and interfaces, always a rich area. What I'd ask you to do is keep an open mind for the kind of sensors that are available.

We just finished talking about algorithms, coming up with smarter systems. And just finally, to reinforce, the best thing to do is keep in mind balancing those three in any practical system. So, thank you very much.