Oticon Technologies to Manage Complex Listening Environments
Recorded November 25, 2019

Presenter: Don Schum, PhD
AudiologyOnline.com Course #33951
Partner: Oticon

- - [Donald] Hello, everybody, this is Don Schum, Vice President of Audiology at Oticon. And what I wanna present today is actually the second show, in a three part series that we have put together on Oticon technologies. The first show in the series was talking about, the technologies that we created to preserve information in the auditory signal. In this second show, I wanna talk about the technologies that we use to allow patients users and users to manage complex listening environments and then in the third show that will be recorded in just a few weeks, we'll talk about the sort of technologies that we use to make sure that we ensure the highest levels of user acceptance. But as I said, today's show is going to focus on how we manage complex environments so let's go ahead and get started. We'll start with the learning objectives and these learning objectives are very important to understand a lot about the way we think about, complex environments and the way patients operate in them.

One of the first things that I think is very important to do is to recognize and accept and deal with the characteristics that define realistic, complex listening environments. Our field has gotten better over the last decade or two about being more realistic about what noise actually is. In the older days when I was coming up as a young audiologist, when we talked about noise, we talked about noises that really did not reflect what real patients run into in terms of problems but we'll talk about what those characteristics or realistic environments are actually like. We'll talk about some of the principles that drive speech understanding, in these challenging environments so that we can put our signal processing routines in context of what really matters to a patient when they're in those environments. And finally, we'll talk specifically about, OpenSound Navigator and some of our other technologies that we use to allow patients to operate well in these environments. To start with, as I like to start with is talking about the why and the why and as far as talking about our approach to signal processing in complex environments is because we firmly believe that this sort of description of what you wanna do, noise bad, speech good, is a gross oversimplification of what true perception is like for a person with normal hearing. We live in a world with other sounds in it. There are sounds in this world that are not just the direct speech that you're listening to. At this moment in time you're probably listening to me in a relatively quiet environment but maybe not maybe you're in a work environment and you're listening to me over headphones in the middle of the day or maybe you're at home and it's perfectly quiet. Or maybe you're at home and you have the rest of your household creating noises. But those other noises are part of your life, are part of your setting, are part of the sound setting. And it's absolutely true that patients with sensory neural hearing loss need an improvement, in the signal-to-noise ratio in order to operate that levels that start to at least approach what a normal hearing person can do, in an environment like this. But that doesn't mean that you need to make all noise go away and you should only listen to speech in perfectly quiet environments in order to understand it. That's an unrealistic goal to have, but it's also a very unnatural goal to have. And so when we approach the idea of trying to create signal processing that allows a patient to operate well in a complex environment, we understand that we need to balance the need to improve the signal and improve the signal-to-noise ratio to allow the person to operate well.

But also maintain a certain sense of naturalness and a certain connection to the environment in order for the patient to have the most natural listening experience possible. And so this really drives the way we go about trying to find solutions that you can offer to your patients. It's absolutely true that when you take a look at what conventional amplification does, the main job of amplification is to improve audibility and to improve the signal-to-noise ratio. Those are two aspects such of Sensorineural hearing loss that drive the ability of patients to operate in complex environments to understand the speech that they wanna listen to. You have to make sure that it's audible to the patient and to the best job as is realistic you need to improve the signal-to-noise ratio. But the devil as they say is in the details and the way you go about that really matters in terms of what the overall impression of amplification is gonna be. Yes, you can improve audibility and yes, you can improve signal-to-noise ratio but you could do in ways that may be better or might not be as good depending on what your approach, what your philosophy of creating signal processing is. At Oticon we have some very specific ideas by which we believe that we can definitely improve audibility and certainly improve signal-to-noise ratio but do it in a way that leaves the patient with a very good listening experience. One of the first things that we need to talk about and this is one of the things that came up in the learning objectives is to understand what are the principles that drive speech understanding on the part of patient? Yes, speech needs to be audible and it also needs to be above the noise level. We know that and the signal-to-noise ratio has to be improved but there's a lot of other aspects to understanding speech in realistic complex environments that will drive how well a patient actually understands speech.

For example, location is an important aspect. Being able to separate out where the speech that you're interested is coming from, from other sounds in the environment is an important cue and it's something that the cognitive system will naturally use. And one of the approaches that we wanna make sure that we do is make sure that we preserve that information as much as possible. What constitutes the noise is an extremely important aspect and the reason why that's so important is that noise is not something that is simple to signal process away. In other words, in the old days, when we used to talk about noise, we used to talk about noise in terms of kind of an engineering description of noise, kind of like narrowband limited white noise and things like that. The reality is noise is whatever the patient, doesn't wanna listen to at that moment in time. And because of that, the characteristics of that noise will determine how easy or how difficult it is to actually remove it from the environment and I'll talk more about that in just a second. Who the talker is matters, some talkers are just more intelligible than others. The support cues that are available how sensitive the patient is to background noise is one of the things, one of the hallmarks of Sensorineural hearing loss is that the distortion component of Sensorineural hearing loss, the usability of the remaining hearing that the patient has varies greatly from patient to patient. And because there is so much variability, the patient's reaction to signal processing will vary significantly, depending on just how much improvement they need in a signal-to-noise ratio. The cognitive system plays a big part of that. Some people are just better at using a minimized set of information than other people are. Some people can just put the pieces together better, they're just better at doing that and because of that, that will very greatly determine how well they can operate in situations where they may not get the full cue set. And finally, one of the factors that we become more and more aware of over the more recent years is the fact that the patient's investment, in the listening task also matters. In other words, patients can dial up or dial down how hard they're working at the task and that's also going to determine how well their total experience is out of amplification. When we talk about complex environments we have to be realistic about what constitutes complex environments.

Like I said a few moments ago, it is basically anything. Noise can be defined as basically, anything the patient doesn't wanna be listening to at that moment in time. So if you're listening to the speech from one person or maybe a group of people who are talking to you, they all might be part of the signal. But if they start talking over each other or they get into a side conversation, then only part of that montage of different voices becomes important to you. There can be movement. In other words, you can be communicating while you're moving or the noise sources in the environment can be moving. If you're standing close to traffic or you're in a restaurant, for example and there's people moving around or something is happening. You can have stable non-speech sources like a blowing of a fan, those typically are the type of noises that are easiest to engineer away or signal process away but they're not, definitely not the totality of the sort of noises that a person with Sensorineural hearing loss has to deal with. There are unstable non-speech sources of sound. One of the good examples is traffic noise that it is changing in time over time. You can see a long-term spectra of traffic noise which tends to be usually under most situations tends to be more low frequency waited but the reality is, is it's changing all the time. And on average over a long period of time it may be low frequency waited but there could be a lot of things going on that's changing as the traffic is actually, moving in the environment. You can have distractions. Sometimes you can just understand speech, better than other situations just simply because there's less distractions going on. There's less cognitive distractions going on, sometimes you're really concentrating versus sort of dividing your attention. There can be less other sensory distractions, going on in the environment and that can drive a person's ability to operate an a complex environment. You can shift focus, for example, like I said, you could be listening to maybe one particular talker for a certain amount of time but then somebody else starts talking and you wanna pay attention to that person. And let's say the person you were just listening to is still having a conversation 'cause they're having a conversation with not just you but maybe a couple other people. But now you wanna shift to be a part of a different part of the conversation and that changes all the time. And you could see some of the challenges in trying to signal process around that. And basically complex environments, you have to think about could be a little bit of everything. It could be a little bit of sitting around with your family making a meal and so you have appliances running, you have conversations, you have side conversations.

Maybe you have a sporting event going on in the TV, in the next room, you have conversations going on in that room, just a lot of different activity that can be going on and patients with Sensorineural hearing loss, need an improvement in the signal in order to understand speech but they don't necessarily wanna be so totally removed from that environment that they don't hear anything else going on. First of all, it's going to be impossible to totally achieve that but second of all, all those other sounds are part of the liveliness of that environment. And you definitely need control over those other sources of sound but that doesn't mean you need them to disappear from the person's life. Think of this situation here. So imagine that the person of interest is that woman kind of in the middle a little bit over to the left in the beige blouse and the brown skirt. And she's, let's imagine that she's the listener in this environment. So she's in a nice, moderately lively cocktail environment with people around, you see a band in the background kind of off behind her to the right a little bit. You see traffic noise happening out there, you have someone running the cappuccino maker over to on the right side over to the edge. You have other conversations going on, you have a fan that's kind of blowing directly in front of her but behind the person she's having a conversation with. And notice that she's having a conversation not with just one person but with three different people. So in this environment, part of the enjoyment of that environment is probably that it is a lively environment. It is a place where she wants to be with other people probably. And she's not having a conversation with just one person within a very narrow angle, right in front of her. She's trying to enjoy a normal sort of social interaction with multiple people. And maybe the music is pleasing. Maybe she likes that little combo that's going on there. And so the idea of those other sounds in the environment that's part of the environment. It's part of the ambiance of that environment, just like the artwork is, just like the architecture is, just like the smells of the food.

Everything is part of that environment and you don't necessarily want to create it in such, put the patient in such a bubble that they don't, they're not allowed to be part of that environment. And so all those factors are things that we think about when we think about what we need to do as far as signal processing. So when we set out to create solutions for listening in complex environments, there are three basic principles that we follow. The first principle is that we wanna make use of the brain. By far the brain is the very best signal processor that we have working for us. And one of our very clear opinions that forms the basis of our brain hearing approach is that speech understanding is a cognitive process, it happens in the brain. There is no speech understanding that happens in the auditory periphery. The auditory periphery's job is to take acoustic signals that are occurring out there in the environment and turn it into a neural code and send that neural code up to the brain. And the brain is where listening and understanding occur. And because of that, we wanna make sure that we are providing the brain the best signal possible. Now we know that we're working at a disadvantage for patients. We are working at the disadvantage that they don't, the neural code that's generated that reflects the sounds in the environment, is an incomplete code, it's a distorted code. It is the sort of partial information that makes it harder on the brain. And we can't change what's gonna happen in the peripheral auditory system. The disorder that has been created because of the hearing loss is what it's going to be. We have to change sound before it enters that system but what we wanna do is the best job possible to change sound before it enters the system so that once it passes through that disordered system it is still in a form that the brain can get the most information as possible from. We can't undo what's gonna happen, in the peripheral auditory system but we can help prevent some of the effects of what's gonna happen. But at the end of the day, the brain can do a tremendous amount of taking information, even if it's partial information, incomplete information, partially distorted information and turn it into meaning for the patient.

And the more, the better we understand how that occurs, the better we can create solutions that try to provide the brain the best possible set of information as we possibly can. So when we've gone about creating, signal processing routines, for speech understanding in complex environments, we always do it with a focus on what is the most important information that we can be feeding the brain? And let's do the best possible job that we can. The second principle that we follow is we wanna make sure that we're preserving the good information that's out there. And in the first show in this series, I talked about cue preservation technologies. And in that show, I talked about the different sort of information that's out there in the world acoustically that the brain can make use of that sometimes gets corrupted by conventional technology. And part of the principles that I talked about in that show is that where sound comes from matters, the bandwidth of speech matters and the details of speech matters. So in that show, I talked about a couple of our different technologies including Spatial Sound and Speech Guard. And those are technologies that we specifically developed to try to preserve as much information in the speech signal as possible and the reason we do that is one is because it's always good to present a very good signal to the brain. But especially in complex listening environments, the better detailed the speech information, the better detailed resolution, about sounds in the environment are, the better off the patient is going to be. So those are very much solutions that we believe have an impact, in difficult listening environments. They are not specifically designed to deal with the noise but they're specifically designed to preserve the sort of information that the brain would use to help understand in noisy environments. So one of the first things that you could do is you could think of it in terms of the old Hippocratic Oath sorta approach. The first thing you wanna do is do no harm and so we have created certain technologies, in our hearing aids specifically to do no harm to the signal in order to try to preserve as much of the natural information as you can. That's not enough. That's not enough for the typical patient with Sensorineural hearing loss but it is a good start. And then what we add on top of that are signal processing routine specifically, one big one called OpenSound Navigator that's designed specifically to try to help improve the signal-to-noise ratio for patients but without creating an unrealistic and an unnatural listening situation.

So the third principle that we follow then is we wanna present the user with a cleaner, rebalanced signal. And what I mean by rebalanced is that there are multiple sources of sound in the environment. And like I said, we don't necessarily, want those other sources of sound in the environment to disappear because those are part of the environment. They're part of that listening environment and there's value in them. There's both an aesthetic value but there's also communication value to those sounds in the environment. But they have to be rebalanced, meaning that they have to be put in a proper perspective so that the patient still gets a good signal-to-noise ratio to operate with but without creating that unnatural listening experience. And if you get too aggressive or if you use signal processing approaches that only focus on doing, maximum signal-to-noise ratio improvement without regard to any other factors, then you can end up with a very unnatural listening experience. So we want to improve the patient's performance but at the same time is preserving a very, natural listening experience for the patient. It gets back to something I talked about a few moments ago and that is what is the true nature of complex sound environments? And when you take a look at the true nature of complex sound environments, you realize that just simply thinking that noise is something that's bad and it's different from speech and we could make it all go away and just preserve the speech. First of all, like I said, that's not a realistic approach and second of all, it's a very unnatural approach. And so, by recognizing what true sound environments are like then that really allows us to get our mind around what we wanna do, when we're creating signal processing routines. If we take a step back before we created OpenSound Navigator and think about what speech in noise technologies looked like back then. We talk about traditional technologies and back in the day, there were two particularly traditional technologies that were being used and are still being used. And one was multi-band adaptive directionality and then the other one was noise reduction. And those were two technologies that we had been using in our field for close to 20 years.

Basically, since the advent almost of digital technologies in hearing aids in the mid 1990s, then one of the first things that happened with digital technologies is that the hearing aid companies started to try to figure out how can we use the power of digital to start improving the sort of signal processing that we could do with hearing aids? And we first developed noise reduction strategies and then adapted multi-band directional technologies and they served us well for a decade or two in terms of doing a good job of taking in amplified signal and trying to improve the signal-to-noise ratios. But there definitely were some limitations. So for example with adaptive multi-band directional technologies, the idea was that you can use the characteristics of where sound is coming from then try to minimize sound coming from the back or sides of the patient and try to focus on sounds coming from the front. And to some degree that those were successful, in improving signal-to-noise ratio but it was very dependent on the specific characteristics of the environment that the patient was in. And the way that those technologies changed, being in a directional to a multi-band directional mode, those kind of varied from company to company about which sort of criteria were used and things like that. But they were all working on a, pretty similar principles about how you set up a directional system to work. When you talk about noise reduction in hearing aids, they also were working under relatively simple principles. One of the basic principles that was being used to drive noise reduction was the observation that speech is a modulated signal. And what that means is that when someone is talking like I'm talking right now, that my energy pattern is going up and down. Every time I produce a phony, it's going to be produced at a different level over time. So vowels, especially certain vowels are, more intense than other vowels. Vowels typically are more intense than consonants. Voice consonants are more intense than voiceless consonants and so constantly as I'm producing phonemes and syllables that my pattern of speech is going up and down, and up and down and up and down in terms of level. And that's been an observation about speech for as long as we've been recording speech.

Think of the VU meter on your audiometer that's bouncing all the time. And the idea being that while as speech bounces all the time, there's a lot of sounds in our environment that don't change in that pattern. Are stable non-speech sources of sound. For example, the blower on a computer or an older computer or on a fan in your kitchen or air conditioning unit that tends to be relatively stable. And traditional noise reduction was basically driven off, the principle that speech changes over time and noise doesn't. So if you find signals that are not changing over time and that pattern that speech typically changes over time, then that's probably noise. If you find signals that are changing typically three to eight times per second which is the modulation rate of typical conversation speech, then you can be pretty sure that that looks like speech. The problem with that approach to noise reduction is that in order to get rid of noise, you have to find it first and you have to define it. And like I said, there was definition of what noise was like in noise reduction and there's definition of what noise is like in traditional directionality which is any SPL coming from the backs or sides of the person. But that typically was not enough to really solve the problem. So if we look specifically in terms of noise reduction, like I said with adaptive directionality, it works pretty well if you're in a situation that's really clean about speech coming from in front and noise coming from the back to side. But there's a lot of situations where those conditions weren't being met and so there was a natural limit on how much adaptive directionality can give the patient, it gives some, it might get three or four even five dB of improvement but they asked about it and those conditions have to be pretty, pretty well defined. When you talk about filtering, noise filtering, for speech versus noise or noise reduction, all the characteristics are driven by the long-term characteristics of the speech versus the noise. Remember, I said that speech is a modulated signal, meaning it changes in level over time but in order to see those changes over time, you have to track it over time. And so you can't have like a moment to moment noise reduction because you have to track to see whether or not you're seeing these three to eight times per second modulations in the speech signal. And so any filtering that was being done with noise reduction tended to be more long-term in terms of the changes.

The other issue with that is that speech is a broadband signal. So even if you can define that the signal is changing, in a modulation pattern over time that looks like speech. The one thing that we know about speech is that speech fills up the bandwidth of a hearing aid and goes beyond the bandwidth of a hearing aid. So you would expect there be modulation, in any frequency region in the hearing aid. If you are playing speech in the background of let's say, a low frequency dominated noise, let's say some sort of machine created noise like a blower or something else like that, that tended to have a low frequency characteristic to it, then you can say, well, in low-to-mid frequencies, I have a relatively poor signal-to-noise ratio. So I can kinda filter out the mid-to-low frequencies to try to reduce the effect of the noise and just amplify the high frequencies. And that's the way classic, noise reduction in hearing aids typically works. The problem is is that a lot of the noises that patients run into like I said, a couple of times, during the show already are other people talking. And with other people talking, first of all, the bandwidth is going to be the same as the speech that you're interested in. And second of all, it's going to be modulating just like the speech signal that you're paying attention to. So the idea that traditional noise reduction, can somehow handle competing voices doesn't make any sense because competing voices have the same characteristics that you're using to define the speech signal that you're trying to save. So there was always a limitation in what traditional directionality could do and what traditional noise reduction could do for a patient. The other issue about them is that because of the way they looked at signals and because of the timeframe over which they had to evaluate signals, they could only change modes in a relatively slow manner. This is the response to a speech in noise signal mixed at a relatively poor signal-to-noise ratio for one of our hearing aids that we had before we released Opn with OpenSound Navigator.

So this is a hearing aid that was using a modern version of conventional directionality noise reduction. And in that series of hearing aids, we actually allow the hearing care professional to set the responsiveness of the system using something called personalization or identities to be either faster or slower in response to changes in the environment. So what you see in this graphic so lively is a more fast response system. Steady is a more slow response system and what do you see in this graphic is you see, two points in the output waveform of the hearing aid. There's a point in which the directionality is starting to kick in. So that's indicated by the red arrow and then you see a point where the noise reduction kicks in after that in time and that's indicated by the yellow arrow. And the timing parameters that were used in these series of products were pretty typical of the timing parameters that were being used at the time across high end products in the marketplace and the timing parameters that are still being used in products that don't have OpenSound Navigator in the marketplace. And what you notice is that even in the very fastest response system so the lively system and directionality, it still takes two or three seconds before you start to see directionality to have its full effect. And in the slower response system it can be up to let's say, five seconds before directionality starts to kick in. And then beyond that, you need noise reduction to continue to do some work. And so you see noise reduction taking another, let's say, four to five to six to seven or eight seconds before it does what it's going to do. But in general, what you're seeing is that traditional noise reduction and directionality, just took time to do what they were going to do. And it was measured in terms of seconds and usually single to two double digit seconds. Like I said, somewhere between three or four or five seconds to maybe up to 10 or 15 seconds before it does something. And that's just because the analysis that was being done by those systems in order to respond to sounds in the environment, was just not able to respond any faster than that. It just needed to know more information, before it could make a decision about whether or not the response of the hearing aid needed to change because of the system. And so that's the, those are the characteristics that conventional directionality and noise reduction at the time. They just were relatively simple systems that were operating in relatively slow time base.

We got to a point where all the manufacturers recognized that that wasn't good enough and we needed to do something different. The problem is that some of the other manufacturers decided that the next best thing that we could do would be to use beamforming. So once we're able to get hearing aids to talk to each other across the head, one of the things that became possible was to create a directional mode that was much more narrow than the standard directional mode that you can create with a single hearing aid. And that's because you have multiple microphones that are separated in space. So a couple of manufacturers decided that since patients were still struggling, in complex environments and that conventional directionality and noise reduction had done everything it could possibly do that the next step would be to create beamforming solutions so very narrow responsiveness of the system. And basically what they're saying is that, well, patients still need more help and there's nothing that we can do to make these competing sounds go away, except create a very narrow listening path for the patient. So the patient can have an improved signal-to-noise ratio, for things that are happening directly in front of them but anything else that's happening in their environment, we're gonna reduce significantly. And that was true that beamforming does what be informing is described to do. Beamforming does improve the signal-to-noise ratio, for things directly in front of the patient. But it also is true that it creates significant attenuation, for other sounds in the environment. To us that wasn't good enough, to us that was creating a very unnatural, listening experience for the patient. We did not wanna move in the direction, we were somehow restricting the patient more than ever. We thought that we can do a better job in terms of creating improvements in the patient's perception but without creating this unnatural listening experience. In a show that I recorded on a couple months back, called "A Deeper Look at Sound Environments" one of the points that I made in that show was that the sounds around us help to create a sense of place.

It's very natural for the human being to want to be immersed in a sound field that brings quality to it, that brings a sense of what that listening environment is. Places are supposed to sound a certain way and if the sound of a place doesn't match what you see or what you expect by being in that place that creates a very unnatural connection or lack of connection basically, between the person and the place that they're at. And so it's a natural human desire to wanna have a more natural sense of the place around you. And beamforming solution simply runs very contrary to that way of thinking. And because of that, we created something called OpenSound Navigator. And OpenSound Navigator was designed, like I said before to create an improvement in performance for patient, in complex environments but without creating an unnatural listening experience, for the patient. Very importantly, OpenSound Navigator was driven off what was originally the Velox platform and now has been improved significantly and updated to the Velox S platform. And there's a lot of numbers on the screen about what the technical specifications of the Velox S platform are. The most important thing to know about the Velox platform is that it runs at a very, very fast speed. And we dedicated the primary improvements and speed in that platform to OpenSound Navigator. One of the major differences between, conventional technologies that handle background noise and OpenSound Navigator is the speed at which it can do things. Like I just said, typical conventional noise reduction and directionality takes several seconds before, it can decide what it wants to do in a complex environment. And that wasn't good enough, that's not the way we wanted the system to be able to work. And so we needed to have a much more powerful engine, in order to run at a speed that we think was necessary in order to really give the patient benefits, in these complex environments. But having a more powerful engine isn't enough to be able to solve the problem and you needed to have some very specific new thinking in terms of the design. So OpenSound Navigator is driven by a very well patent protected approach to how we want to handle complex environments. And this is where our approach to signal processing really shows up.

There are a few very important differences to the way OpenSound Navigator works compared to conventional directionality and noise reduction and compared to beamforming. The first very important difference is that there is an analysis function that happens in the product before, any signal processing is allowed to be applied. And this is important because in traditional hearing aids, the way they're wired, the sound comes in and the first thing it does, is it goes into a directional mode or it goes into a directional system. And directionality cleans up the signal-to-noise ratio as best as it can and then the signal is fed forward to a noise reduction system that does whatever it can but the noise reduction system doesn't know what the directional system knows, it just gets a signal and it's basically dumb to anything that happen. And then once it's gone through the noise remover section, then it goes on and gets amplified and compressed and whatever it does. And the, one of the defining unique characteristics of OpenSound Navigator is this idea that we have an analysis phase up front. And in that analysis phase a lot of characteristics about the sounds that are coming into the hearing aid in terms of location and other details that I'll talk about in a moment, are all specified up front and that information is saved. And it is that information about the environment that is used very much to allow much more complex reaction to the system over time. The second very unique characteristics of OpenSound Navigator is that is part of the analysis phase, it uses the directional microphones on the hearing aid to simultaneously create two different polar plots to describe the environment. So again, nothing as far as signal processing has happened so far, in terms of changing the system, this is still part of the analysis of the environment. And by taking a look at, a 360 degree look at the environment and comparing that to a backwards facing polar plot, that allows the system to in a way that had never been used before and still hasn't been used yet in hearing aids to do a much better description of exactly what the sound environment is. If you can take a 360 degree look at the environment and directly on an ongoing way compared to a backwards facing look at the environment, you have much better definition of what's in front of the patient and what's important to that patient. This isn't the only analysis that it does to the environment but it's a very important part of the way we define what's happening around the patient. Once that information is stored, then decisions can start being made on how to change the sounds in the environment.

And that's where our directional system, what we call our balance program starts to kick in. So it starts making very selective, changes in the responsiveness of the hearing aid in different frequency regions, pointed in different directions to the back and sides of the patient. And at that moment in time, you are using a different approach directionality that has been used in conventional directionality. But it's a very precise and very rapidly updated system that allows you to very selectively, pick out sounds in the environment and to attend to them. And then the third extremely unique approach that is defined in OpenSound Navigator is the fact that this analysis information is also fed forward to the noise removal system. And that is very different than anything that's happened in the past because in the past, like I said, noise reduction typically didn't have any clue about what the directionality system that happened right before it did to the signal. But with OpenSound Navigator because we do the analysis upfront and because we store that information. And because we share it with the noise removal system, then the noise removal system is what has been termed spatially aware. Meaning that not only does it know the types of signals that are in the environment, whether the speech or non-speech but it also knows where they're coming from. And that's part of the power of that system is that it really has a much better look at what it wants to get rid of and what it wants to preserve. With all this signal processing going on and all this analysis going on, there's a lot of different things that we can do to the signal. When we capture the data in this analysis system, we are capturing it in multiple channels, 16 channels. We're capturing level, location, speech versus non-speech and other characteristics of the signal. And by capturing that data and storing it and using it and updating it on a regular basis, it allows a very rapid and very targeted application of the principles of directionality and the principles of noise removal to create that improved signal-to-noise ratio for the patient but without creating an unnatural listening experience. When we, a few moments ago I talked about how we needed a more powerful engine to run this.

And because we have that more powerful engine, the Velox S platform, we are able to update the reaction of the balance system or the directional behavior of the product and the noise removal system over 100 times per second. Meaning every 10 milliseconds or less we can update the characteristics of what sounds in the environment is trying to minimize. And like I said, the the minimal timing that you will see in conventional directionality or noise reduction is maybe a three second update and oftentimes it can be much longer. But this graphic just is created just to remind you what the difference between 10 milliseconds and three seconds looks like. So you can have updates that are happening, compared to what you can get in a traditional noise reduction or directional system, at least 300 times faster, if not 1,000 times faster, depending on the update speed of the system in OpenSound Navigator, than you have in traditional noise reduction and directionality. And I'll give you some more evidence of how that actually shows up in actual performance of hearing aids in a little bit later in the show. But the idea being that the engine of Velox S allows us to run at these speeds. To give you a little bit more details on how to think about OpenSound Navigator, think about the balance system or the directional behavior of the product. Think of it in terms of gain levels for very specific locations, in the environment for very specific types of sounds. So in this case, if you have speech coming from in front of the patient, then what you're going to see is gain applied at different levels around the patient as it identifies sources of sounds in different frequency in different locations and environment, in different frequency regions and with different characteristics that make it either a speech or a non-speech signal. And because of that, we can more precisely adjust how it's gonna handle the levels, coming from these different sources of sound.

So think of it much more in terms of attacking individual sources of sound as opposed to thinking about it as attacking certain directions from around the patient. And because we can update, we can handle things like movement in the environment, we can handle things like sounds starting and stopping or sounds being erratic in terms of how they change over time, some of those non-stable sources of sound that I talked about. But again, think of it as very precise and targeted noise part or gain controls for different sounds, coming from the environment. In general, with OpenSound Navigator we try to create an Opn field, definitely in the front half of the listening plane for the patient. In most naturalistic communication environments, if you're having a conversation with more than one person. For example, think of that photo I showed you a while ago with the woman being in that art gallery opening or whatever that was. She was having a conversation with three other people but they were basically in the front half of her listening field. So what we try to do is try to keep sound relatively well preserved in that front half of listening field, especially well defined voices. As the sound moves to the behind the patient, we're are more likely to try to reduce it but the more it looks like a well defined speech signal, the less likely it's going to be reduced.

On top of that, then the noise removal system can go ahead and attack other sources of noise in the environment, even sources of noise that are in front of the patient. And again, it's because of the analysis function and the way the systems interact with each other. And so what we try to do is create a relatively Opn sound field in front of the patient with some reduction of other sources of sound that might even be in front of the patient. But if we're going to do significant reductions, it's gonna tend to be behind the patient. And what that allows the patient to have is a listening field that basically falls in line with what they see around them. That they had that more natural sort of realistic sort of connection between the sounds that they're hearing and what the environment looks like to them. Again, it's a more naturalistic way to deal with the environment. And again, it's not that we don't do anything to the sound in the environment. That's a misnomer to think that OpenSound Navigator just allows all sound to come into the patient because that would be overwhelming to the patient.

But what we tend to do is keep a much more open listening field but allow very aggressive and very specific attacking of sounds that would be competing with the voice that you want to present to the patient. And again, it's that rebalancing of the sound environment. Like I said, we don't wanna have a signal-to-noise ratio that's infinitely good because that's unnatural in the environment. But we want to rebalance the different sources of sound in the environment so that the patient has a working opportunity to hear and understand speech well, without being lost in all the sounds in the environment. And again, it's driven very much by the way the OpenSound Navigator has been designed with the unique features of being able to have this analysis function up front, this comparison of two different polar plots as part of that analysis. And then this idea of giving the noise removal system information about where the different sources of competition are gonna be occurring, from around the patient. So that's a lot about how the system is designed. It's important to understand what makes it different than traditional systems, in the marketplace. So if you compare OpenSound Navigator to traditional directionality of noise reduction, the things that are basically different compared to those systems in is the fact that we have the analysis phase with OpenSound Navigator, we have spatially-aware noise reduction, like I said is very important. The criteria for noise reduction changes because we're working at a much greater speed and we have much better information about what the sounds in the environment are. It's not just using long acting, modulation-based characteristics of the speech signal to tell us what the speech signal looks like. And then very importantly, the update speed is much, much faster with OpenSound Navigator compared traditional systems. Compared to beamforming, there are some differences too. Again, we now have the analysis phase as part of OpenSound Navigator which is not something that occurs in beamforming systems in the marketplace. Again, the update speed is very different, between the systems and I'll show you that in a few moments. And very importantly is this protection of clearly defined speech sources, especially in the front half of listening field but even very clearly defined speech sources, coming from behind the patient.

Whereas beamforming tends to be very aggressive about reducing everything, except the sounds that are in a very narrow listening window, directly in front of the patient. That's well and good in terms of what we set out to create but what I wanna do now in the last section of the show is present to you the evidence that we've collected over the last few years to show the effectiveness of OpenSound Navigator. So the first work I wanna talk about is some acoustic that, recordings that were made with KEMAR. And so these were KEMAR recordings, comparing the Opn 1 product to a beamformer and to a couple different products with traditional directionality and noise reduction. The input signal was Speech in Speech-shaped noise and most of the recordings were done at a signal-to-noise ratio of plus five dB signal-to-noise ratio. And we're looking at, in these recordings is a technical analysis of the effect on the signal-to-noise ratio. So this is not patient performance but this is just showing you what the products actually do over time. As far as the traditional directionality and noise reduction systems, those are very modern hearing aids. So the two products that we use that were using traditional directionality and noise reduction are high end products, currently in the marketplace that define part of the high end market in the marketplace but they are using traditional directionality and noise reduction as part of their way of dealing with complex listening environments. This is the automatic configuration. So it was a classic sort of flattish, modern hearing loss and this is the first listening condition that we use. And I'm not sure why that one loud speaker up to the back left is black, it should be white. So in this recording setup, just to orient you, the KEMAR was wearing a hearing aid over on his left ear and he was facing forward on the screen or down on the screen. So speech was coming from in front and speech-shaped noise was coming from behind him, from the two back angles coming from behind him. And so this is kind of a softball condition for modern hearing aids. A modern hearing a should be able to do a good job in an environment like this. So let's see what the effects look like. So in order to take a look at what it's gonna look like, these recordings were 35 seconds long and they included three different sections. So in the first section that goes from about here to here, there was speech coming from one talker and it was mixed, like I said at a plus five dB signal-to-noise ratio. Then a second talker started for about about eight seconds or so and then the first talker and the second talker were mixed for the last section. So it's talker one, talker two and then talker one plus two and that's what you see.

This is the output out of the Opn hearing aids. So what you notice was that relatively quickly within the first second or so, there's no speech here so there's not going to be a super aggressive response, in the hearing aid but it does eventually get rid of the noise in the environment and then becomes very stable. And what you notice is once the speech starts, the background noise becomes relatively stable over time. There's very little fluctuation in the background noise and what you see is a nice improvement, in the signal-to-noise ratio. And what you wanna do is take a look at the peak ratio. Up here this is the unaided signal coming in and this is the aided signal coming out here. And so what you're seeing on the screen, I know you can't read the y-axis up there but you're seeing a good six dB plus improvement, in signal-to-noise ratio unaided to aided. And the way you wanna judge this is basically take a look at the difference between the ongoing background noise which would be the lower pink line so this one right there and the typical peak level that's coming out of the system. And if you compare this difference between the two lines so this difference, that's where you get about a five to six dB improvement in the signal-to-noise ratio, between those two systems. This panel will give you the comparison between Opn 1 and then product A and product B. In product A and product B we're using, traditional noise reduction and directionality to handle the environment. And basically what you wanna take look at is how quickly the system seemed to respond to the speech plus noise situation, okay? So if you take a look at it, like I said within the first couple seconds, even if it was just noise coming from behind, OpenSound Navigator is still gonna realize that there's absolutely no speech that it wants to do anything with and reduce it and reduces it further quickly and keeps it down. But if you look at product B and product A and especially in product A, what do you notice is that it takes that product a good almost 20 seconds, before it can finally decide what it wants to do with that background noise situation. And remember, these recordings were done with speech coming from directly in front of KEMAR and noise coming from behind KEMAR. So any product that has directionality and noise reduction should be able to handle that system pretty effectively, that setup pretty effectively. But in product A and even to some degree in product B, it takes many seconds before the systems can identify and respond to that noise in that environment. And that is a reflection of the signal processing routines that we put into OpenSound Navigator and very importantly the speed at which they can analyze and update the setting of the hearing aids. And that shows up very clearly on this one. So even though that these are modern hearing aids that Opn is being compared to, these are modern hearing aids, using more traditional approaches to directionality and noise reduction and you see the differences show up in terms of the responsiveness of these hearing aids in this environment. And like I said, it was pretty much of a softball sort of environment to make recordings in. To compare Opn to a beamformer situation, we set up this recording situation. So in this case, you have the speech coming from talker one, one second, okay, the speech coming from talker one, coming from directly in front of KEMAR and then speech from talker two coming from 90 degrees off to KEMAR's left. So it's on the side that we made the recording but it's off to the side 90 degrees. And then speech-shaped noise is coming from the side and back of KEMAR at that point in time. And so this was a situation that we set up specifically to be able to compare what OpenSound Navigator would do, versus beamformer in this sort of environment. And this is what the output looks like. So here you see it with, you see the unaided signal and like I said it's talker one coming from the front, talker two coming from the side and then talker one and two coming from both the front and side. And this is the output from Opn. And what you notice is that with the speaker coming from in front you get that good nice clear signal-to-noise ratio coming.

When the speaker shifts over to the side and remember that speaker normally, was a little bit lower level, if you could notice on the input level. And so although the signal-to-noise ratio isn't as dramatic as it is when speakers come in front, it's still relatively well preserved. And the idea being that the sensitivity pattern of OpenSound Navigator is set up so that you still preserve speech coming from directly to the side. And remember, like I said, we wanna keep the front half of the listening field open so that if someone, if you're having a conversation with more than one person and there's someone a little off to the side of the conversation that you don't have to turn your head and stare directly at their face like you would with a beamformer system in order to get a signal-to-noise ratio improvement for that talker. If we now move to comparing Opn to a beamformer and this is product C which is using a beamform approach. Notice that when the speech is coming from 90 degrees off to the side that the signal-to-noise ratio, the peak sitting up above, the noise level in the environment, is significantly better for Opn, than it is in the beamformer situation. Now the beamformer is doing exactly what a beamformer is supposed to do. So the beamformer is working very well as a beamformer but you notice what ends up happening is that that creates a very, a more narrow listening field, for the listener which is a beamformer's job to do that. But like I said, we just consider that to be an unnatural way of processing the signal for the patient. In a second series of projects that we did, we wanted to see how patients performance was comparing OpenSound Navigator to traditional directionality and to beamforming. So in this study what we did is we had a listener here in the center and they were listening to speech coming either from directly in front of them or to speech coming from 45 degrees from one angle or the other. And where the speech was coming from, was randomized on each trial. So the patient didn't know where the speech was coming from and the idea was how well OpenSound Navigator can handle this relatively, complex but kind of naturalistic situation compared to traditional directionality and beamforming.

This is the data when you compare, when you're talking about only listening, for the person coming directly in front. So these are only the trials and where speech was coming directly from in front of the listener. And what you notice is that there is a significant improvement, if you compare OpenSound Navigator to traditional directionality and noise reduction, you see a significant improvement. But you see no difference between the performance of OpenSound Navigator and beamforming in that environment. You would expect beamforming to do well in this environment, that's exactly what it's designed to do. Is allow you to hear speech coming directly from in front of you well but the important part to notice in this data is that OpenSound Navigator could also do a good job of improving the signal-to-noise ratio, compared to beamforming. So OpenSound Navigator can do what a beamformer does and it does it well and it matches it in terms of performance. And both the beamformer and OpenSound Navigator clearly outperform, traditional directionality in that environment. But the point being is that there's no additional benefit of going to beamforming. Now if we move to environments in which the speech could come from either plus or minus 45 degrees, then things get very interesting because now you see a significant improvement comparing, beamforming to OpenSound Navigator in that environment. So what you see is not only can we match the performance of beamforming when speech is coming directly from in front but we can also significantly outperform a beamformer, for speech coming from other angles. And again, that's exactly that sort of cocktail party situation that I showed you a picture of early on. That's the situation where we wanna create a natural link between the patient and the environment by being able to show how well the system can work. And of course, we significantly outperform, traditional directionality in that environment also. Overall, if you combine the data across all three speakers, you see a steady improvement, going from traditional directionality to beamforming to what OpenSound Navigator can provide to the patient. The final set of data that I wanna talk about is to talk about the implications of improving performance in noise. And so in this project, we're comparing OpenSound Navigator on versus off. And we're measuring performance in two ways. And we tested speech in noise across a variety of different signal-to-noise ratios, from about negative 12 all the way up to about plus 15.

And we test it like I said with OpenSound Navigator on and off. The first set of materials that we used was designed to give a classic speech in noise, performance percent correct. So by testing at the different signal-to-noise ratios, you see the classic sort of improvement over time as the signal-to-noise ratio gets better, you see performance goes up. The solid line is for OpenSound Navigator off and the dotted line is for OpenSound Navigator on. And what you notice is that we can improve the signal-to-noise ratio performance at the different signal-to-noise ratios significantly as you move across time. And typically, this is equivalent to an improvement of at least about six dB signal-to-noise ratio improvement, equivalent to an improvement of about five to six dB signal-to-noise ratio at least. And again, it happens across a variety of different signal-to-noise ratios but you see a significant improvement to the left on an input-output function for, performance intensity function. And that's exactly what you expect from a noise reduction system that across different signal-to-noise ratios, it can improve patient's performance, until the signal-to-noise ratio gets very good. But like I said, this is about the implications of improving the patient's performance, in noisy environments. So the way that we took a look at the implication to this is to measure the listening effort that the patient had to put in. And instead of measuring it subjectively where we asked the patient to rate it, we measured objectively using pupillometry. Pupillometry is a measure of the size of the pupil and it's been well established in the scientific literature that the more cognitive effort you put into the task, the more open the pupil is. In other words, you get a wider opening, in the pupil when the person is putting effort, into a listening task or any other kind of cognitive task. And that's measured by a measurement of pupil dilation and what we noticed is that with OpenSound Navigator off, when the signal-to-noise ratio is very good, you had a baseline pupil dilation. And as the signal-to-noise ratio started to get worse, you saw an increase in the pupil dilation.

In other words, the pupil is opening because the person is putting, more concentration in the task. But very importantly, one of the things that we noticed is that once this signal-to-noise ratio got, very, very poor, the person basically stopped trying because pupil dilation started going back down. Meaning that there's a point at which humans just simply give up and that's pretty easy to understand. That if you're trying to do a task and you're not getting anything right on the task, you basically stop trying, you just guess and you just go on with that. So you see a very interesting pattern where you see a baseline when listening conditions are good. You see some increase in listening effort, when the situation gets more difficult and then you see a recovery phrase where the person just stops trying. That solid line is with OpenSound Navigator off and then with OpenSound Navigator on, you see a very parallel set of performance but with a couple very big differences. First of all with OpenSound Navigator on, you see the mean pupil dilation has dropped significantly. You see the same sort of pattern where it starts at the baseline and as the signal-to-noise ratio starts getting harder, pupil dilation starts to increase. But what you notice is that the give up point now is at a much poor signal-to-noise ratio. In other words, the person is, because they're getting more speech information and they're understanding speech information better, they're willing to try harder and they're willing to try into more difficult signal-to-noise ratios. So we weren't seeing a give up on this task, until the signal-to-noise ratio got close to negative 10 db. Meaning that we allowed patients in, the OpenSound Navigator allowed patients in this task to be willing to put in the listening effort 'cause they're getting some payment or some payback at a poor signal-to-noise ratio. And again, the idea being that we want to allow patients to be operating in environments that are more challenging. And by being able to improve the signal-to-noise ratio or improve performance across our brain's, just signal-to-noise ratios we're able to demonstrate that they're willing to give it their all in a broader range of environments. Overall, when you take a look at these three different data sets and then take a look at the performance of OpenSound Navigator in three different ways, you can see how we have been able to do a very good job of improving performance but at the same time trying to maintain that natural listening experience, for the patient and that's exactly what we wanna do. Again, this is based on our brain hearing way of thinking, we wanna make sure that we are feeding the brain, the very best information. We have to preserve good information.

Like I said, we do that but then we also need to provide the patient a benefit, in those complex environments. But to provide that benefit without creating, an unnatural listening situation. So we go back to our friend who wants to be in that environment and be interacting with other people and to enjoy the setting that's exactly why we created OpenSound Navigator. It's to allow that person to have that very natural interaction with that environment but to have the performance improvement that is necessary to allow the patient to perform well in an environment such as that. We are very happy with the way that we have moved the needle on what it takes to do signal processing in complex listening environments. OpenSound Navigator has been around for a few years now. We've been able to improve it, we've been able to improve the platform, we've been able to improve the products that we put it in but fundamentally, the ideas behind OpenSound Navigator, are still there and they're still there because they work extremely well for patients. As always, if you have any questions or comments on this or any information you can contact me, my email is on the screen. And with that, I wanna thank you for your time.