ANDREWSEYBOLD

Andrew Seybold, Inc., 315 Meigs Road, A-267, Santa Barbara, CA 93109
805-898-2460 voice, 805-898-2466 fax, www.andrewseybold.com

Wireless Noise

October, 2012

Prepared by:
Andrew M. Seybold
CEO and Principal Consultant
www.andrewseybold.com

WIRELESS NOISE

Executive Summary

Commercial wireless networks (cellular networks), or in fact any wireless networks, are optimized to provide the best performance possible for the users. The networks are designed to offer as much coverage and capacity as possible so they can serve as many users as possible. The engineering involved in the design and deployment of these networks is part science and part art, and a lot of time is spent optimizing them after they are deployed.

One factor that cannot be engineered out of a wireless network is the impact of two different types of noise in these systems. The first type of noise, with which we are all familiar, is the noise picked up by the microphone during a call. Even when the user is not speaking, background noise is being sent to the person on the other end of the call. Background noise has a detrimental effect on the way a voice call is heard at the other end of the conversation. The more noise there is, the more difficult it is for the person on the other end of the call to hear the voice and understand what is being said. This has been a problem since two-way radios were introduced in the 1930s and has become a major complaint of those using today's commercial wireless networks.

We take for granted the fact that we can make and receive voice calls almost anywhere we are today. Commercial network operators have spent $billions in building their networks as well as interconnecting them with both landline telephone systems and other wireless networks. Still, making a call in a noisy environment can be a real problem. Many times the calling party will raise his/her voice in an attempt to overcome the background noise and turn up the receiver volume in order to be able to communicate in noise filled environments.

There have been many attempts at filtering out background noise to render voice conversations more intelligible. However, until recently, most of these attempts have been only marginally successful. This is because background noise comes from multiple sources. People talking near the caller, announcement speakers in airports, cars and trucks passing on a road, and background music are all different types of noise and until now, most attempts to filter them out have been with systems that identify specific frequency ranges and filter out some of the noise that falls within them. Sometimes this also has the adverse effect of filtering out some of the caller's voice frequencies as well.

Voice quality improvements deliver advantages for both wireless customers and the network operators. The benefit for the network operator is that implementing noise suppression may result in increased network capacity. When intelligent noise suppression is enabled, significantly reducing background noise during a voice call, each device in use within a cell sector operates at a lower data rate, thus more devices have access to the voice capacity of a given cell sector. This is important especially in metro areas where there is a greater population of customers requesting service within a single cell sector. Invoking noise suppression may result in more usable capacity within the cell sector. This is true even when network operators deploy wideband voice services and many of them are implementing these services today to improve voice quality.

One company, Audience, has invented a much better way to filter out background noise. It doesn't simply build a filter to mask noise; it actually analyzes the noise, measures where it is in the audio band, and then in real time, applies the proper amount of filtering to maximize the caller's voice while minimizing the noise. The system is so good that even if the background noises change, for example, a police siren is heard; the system recognizes the new background sound and filters it out as well. In a recent demonstration of the system, I was on a call where the speaker was in a room with loud music playing. Without the Audience noise cancellation it was very difficult to understand the voice, however when Audience's noise cancellation technology was invoked, the music was removed so that understanding what the caller was saying was a non-issue. Many handset manufacturers and network operators have recognized that Audience has a superior method analyzing voice versus noise, and more are soon to follow.

The advantages of Audience technology do not stop with removing background noise from a voice call. The other type of noise that impacts wireless network operations is generated within the radio bands, resulting in noise that reduces the range of a system and may cause customers to experience dropped calls or having their call blocked. Audience provides additional noise suppression that can reduce this type of noise interference as well, thereby providing a better user experience and/or saving network operators money by helping to reduce this type of noise interference.

I do not know of any other company that has addressed the reduction of noise in this manner. Audience has done extensive testing and has developed a system that works and works well. It has been built into commercially available devices for a few years now, and it has proven to be effective over time. It is clear that Audience's research and technology advancements are setting the bar for all of the industry.

Table of Contents

WIRELESS NOISE

Introduction

Noise is wireless customers' and network operators' worst enemy. There are two different types of noise that impact wireless performance from the customers' perspective. The first is background noise when on a call, which makes it more difficult for the listening party to hear. Conversely, having to listen to an incoming voice call that has a lot of background noise can be both annoying and, in severe cases, make it impossible to hear the words being spoken.

The second type of noise is the noise generated by the radio system (cell system) that can affect the reception of a signal both on the device and at the cell-site end of a conversation. Every portion of the spectrum has a certain amount of noise within it, there is no such thing as clean radio spectrum. When there is also interference or when a received signal is very weak, the voice or data clarity can be deteriorated and the distance over which a handset will work can be shortened by as much as 20% or more. (In technical terms this is known as the SNR or Signal to Noise Ratio. The higher the signal and the lower the noise, the more clarity in the conversation.)

Both types of noise can impact the usefulness of wireless devices so network operators are constantly striving to find ways to reduce both types of noise. Reducing the noise transmitted along with speech renders the speech much more intelligible even when the device is in a harsh environment, and effective suppression of unwanted background noise will enhance the overall perception of performance of the network. As an added benefit, the talking parties will not have to raise their voices to drown out the noise around them. Because we tend to use our wireless devices anywhere, background noise can be a significant factor much of the time. Further, with some types of variable rate vocoders, if the background noise is reduced then the data rate used to carry the voice over the system is reduced permitting more calls in a given cell sector.

The second type of noise is the electrical noise generated on the portion of the spectrum being used for the voice call or data session, which can also impact the quality of the call. If the electrical noise at the cell site is too high, it can also limit the effective range of a wireless device. Network operators are constantly fighting the battle of noise within the spectrum on which they operate. This interference or electrical noise can be caused by nature or by other transmitters at the same location that are in close proximity.

Both types of noise are disruptive to wireless network customers so network operators and handset vendors are constantly examining ways to reduce these two types of noise and thereby reduce the problems they create. Noise of both types will continue to worsen and will plague us on a daily basis. Electrical noise at a cell site will worsen because as more antennas and radios are installed at each site, the more the noise level in a given portion of the spectrum increases. Noise that interferes with our voice and data conversations will also continue to increase as we use our wireless devices more frequently regardless of our location.

The good news for all of us is that there are many smart people working for companies large and small who are tackling one or both of these noise issues. While noise will never be 100% mitigated, it can be

reduced to a point where the wireless devices we use will operate more reliably and our user experience for speech and data will be more consistent.

The commercial wireless space is currently transitioning from historic analog narrow-bandwidth voice to what is being called wideband voice services. Wideband voice increases the amount of bandwidth allocated to each call thereby enhancing the audio quality of the call and audio in general. Music, for example, sounds better on a device using wideband audio.

The downside of this move to wideband audio is that, during the transition, customers will be moved back and forth between narrowband and wideband voice depending on where they are within the network. Work is being done to minimize the differences between narrowband and wideband voice services so customers will be less aware of when they are moving from one to the other. While wideband voice provides better audio quality, it also introduces additional background noise issues since the audio from a device occupies more bandwidth and uses a higher data rate.

Voice Background Noise

One of the most annoying issues for wireless voice users is that on many calls the speaker's voice is lost or partially covered by background noise. One case in point, I just finished a conference call with ten participants. One party to the call was at home, his children were playing in the background, and his dog was obviously joining in. Another party to the call was in a car with its heater running and lots of road noise, while another was in a train traveling up the east coast and train noise and conversations around him could be clearly heard. The result was that many of the participants had to repeat themselves. We tried having all of us mute our phones but several of those on the call did not have any idea how to accomplish that. As a result, the call was longer than it needed to be and the distractions with the background noise detracted from our ability to focus on what was being said.

Most wireless phones have some type of noise suppression built into them, but most phones on the market today have a single microphone and noise suppression is fixed by the software. Thus different types of noise are not compensated for and the result is that those on the receiving end have to struggle to hear the voice of the speaker. Further, there is a natural tendency to speak louder when responding because you know the party on the other end is in a high-noise environment so you speak louder to overcome the noise.

For years, advances have been made in noise suppression technologies. More recently, many devices are shipped with two microphones built into them. Having two microphones enables noise suppression software to function better since it can be written to compare the audio from the two microphones and cancel out some, but perhaps not all, of the background noise. But here again, most noise suppression software is set up to filter out some of the background noise. The problem is that background noise comes and goes, changes in frequency (e.g., baby crying to a car horn, to a train whistle) and each of these background noises has different characteristics. The result is that while most wireless devices have some form of noise cancellation built into them, it is not nearly as effective as it could be.

Noise becomes even more of an issue when voice calls are made on tablet devices. Typically, when a tablet is used for voice, the person is some distance from the device, therefore the voice conversation is more susceptible to background interference. Some of the newer tablets, in an effort to minimize the

background noise issue, are being designed with two microphones. This permits better control over the background noise that is sent over a wireless system.

Recently, the wireless device market has seen a totally new and revolutionary form of built-in noise suppression. These devices include the Samsung Galaxy S II and S III, Galaxy Note, Captivate Glide, DoubleTime, Focus, and Rugby Smart, the HTC Titan and Vivid smartphones and Jetstream Tablet, the Pantech Pocket and Breeze III, the Xolo X900 from Lava in India, the first smartphone to launch based on the Intel Medfield processor, the Huawei Ascend D1 Quad XL and Xiaomi Mi2 in China, and the recently announced Motorola RAZR i. While we see proliferation of this technology increasing in EMEA and APAC territories, it is apparent that the leaders in adoption of this new voice suppression technology are AT&T Wireless and Samsung here in the United States

As can be seen, there are a number of factors that need to be taken into consideration when designing noise suppression systems. How much and what type of background noise is there? Is it at the low end, middle, and/or high end of the audio range? Is it in the same audio range as the spoken voice? Is the device and network making use of narrowband voice only or a combination of wideband and narrowband voice? Does the device have a single microphone or does it employ two microphones? How does the background noise change over time and how well can a noise suppression system handle sudden new sounds such as a siren in the background?

All of these variables need to be taken into account and since many of these change during the course of a conversation, using a fixed form of noise suppression technology does not provide a consistent user experience. Further, a number of these variables will change during a normal wireless voice conversation, thus the noise suppression chosen for a specific device or by a network operator must be capable of changing characteristics on the fly. Audio delivery is a complex issue and must be handled differently within the wireless world than in the wired world to which we are accustomed. As we move from analog voice over our wireless systems to Voice over IP, and from narrowband to wideband voice capabilities, the typical user device will have to adapt to all of these variables if the user experience is to be both consistent and pleasing to the ear.

The Technology

Audience's technology is a chip, branded earSmart, that is built into the phone combined with a lot of very sophisticated software that provides the best noise suppression I have ever heard. The secret of this technology is not only the end product but also the research that has gone into building the chip and the algorithms that make it work. The company has invested heavily in voice and audio research including devising a set of tests that are conducted in a special sound chamber. A dummy or living person is outfitted with microphones placed at strategic locations near the ears and then the dummy, or person, is bombarded with many different types of real-world sounds. However, more importantly to me is the fact that these tests are also run in real-world environments and under a wide variety of noise conditions. The information from this data is then used by the chip and software engineers to develop smart algorithms that are constantly monitoring the audio levels via the phone's or tablet's microphones, so that the proper amount of filtering is applied to remove as much of the background noise as possible while protecting the remaining voice quality. The demonstrations I have witnessed are difficult to comprehend at first but the results speak for themselves. In one of the tests I was on a phone with a caller who was in a room with music blaring (not playing, blaring) from across the table. The conversation on the phone was unintelligible with the noise suppression system turned off, but when it

was turned on, the background noise in the form of the blaring music was reduced to a very mild level. The voice level of the person speaking was distinctive and each word could be clearly heard.

After I had a chance to sit down with the architects of this system and have the technology behind it explained to me in detail, it made perfect sense. Below (Figure 1) is one of the charts they produced that shows the performance of this system with various types of background noise. The background noise runs the gamut of pink noise, car noise, street noise, babble (conversations), voice, music, and others. In this case, you see the difference between the signal to noise ratio (SNR) with the system disabled and when it is enabled. Further, the louder the background noise, the more difficult it becomes to filter out without adversely affecting the speech. In reality, the trade-offs between voice comprehension and the amount of background noise that can be filtered out are many. Applying too much noise suppression can have a negative impact on the intelligibility of the voice, but not filtering enough background noise can also cause a conversation to become garbled and unintelligible.

Below are a series of graphs provided by Audience that break out the different types of background noise so we can see exactly how much difference there is between having the Audience technology invoked and the level of noise when it is not.

The following figures show how noise reduction varies with the amount and type of noise. The basic measure of the amount of noise relative to speech is Signal to Noise Ratio, or SNR, expressed in decibels (dB). As a noise suppressor reduces noise while leaving speech intact, the amount of improvement provided by such a system is measured in terms of Signal to Noise Ratio Improvement, or SNRI, also measured in dB. The SNRI metric has been standardized in the ITU-T as part of Recommendation G.160. The figures below show the amount of improvement, SNRI, as a function of the amount of input noise, measured in terms of SNR. Each colored curve is for one of six different types of noise, ranging from very steady, or stationary pink noise, to highly variable, or non-stationary, single voice and music. Most noise suppressors have been designed for stationary noise types. Recent advances in noise suppression technology allow very good reduction of highly non-stationary noise while maintaining speech quality.

These results are based on the EVRC speech codec (Enhanced Variable Radio Codec). In phones in the field, the EVRC speech codec incorporates a stationary noise suppressor, referred to here as IS-127. Using fixed-point bit-exact reference code, it is possible to examine the performance in noise when the IS-127 suppressor is disabled, as a baseline. Figure 1 shows the results, SNRI vs. SNR for six noise types, for the EVRC codec without the intrinsic IS-127 noise suppressor.

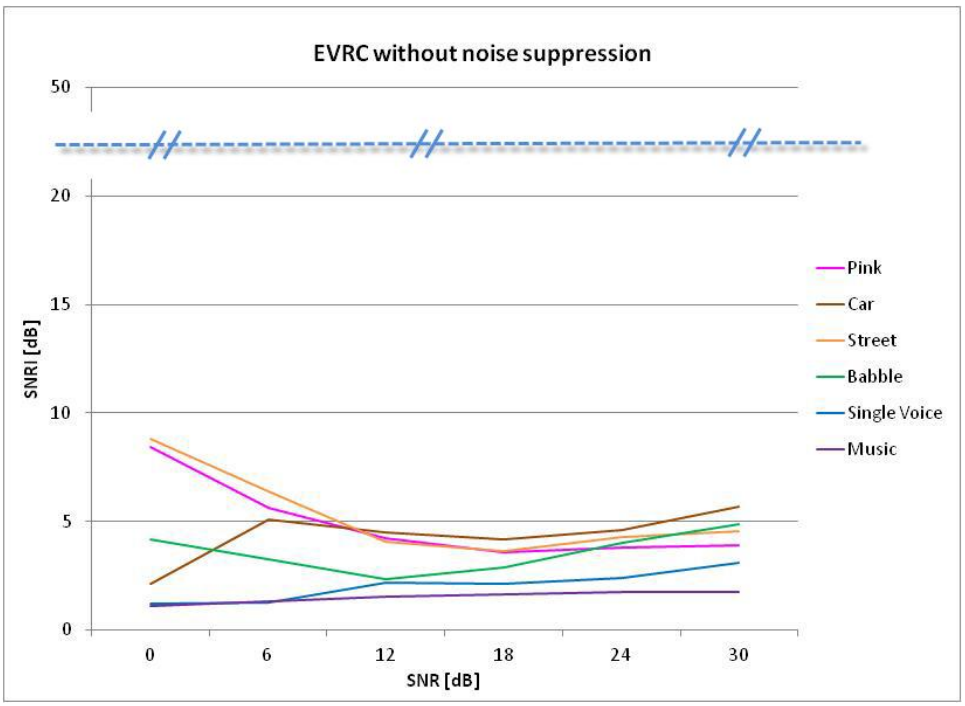

Figure 1. EVRC without noise suppression

Note that for all noise levels (SNR) and all noise types, there is relatively little suppression. The small amount of suppression is due to the inherent bandpass filtering in the speech codec.

The next figure (Figure 2) shows the results for the same conditions, but for EVRC with the intrinsic noise suppressor enabled.

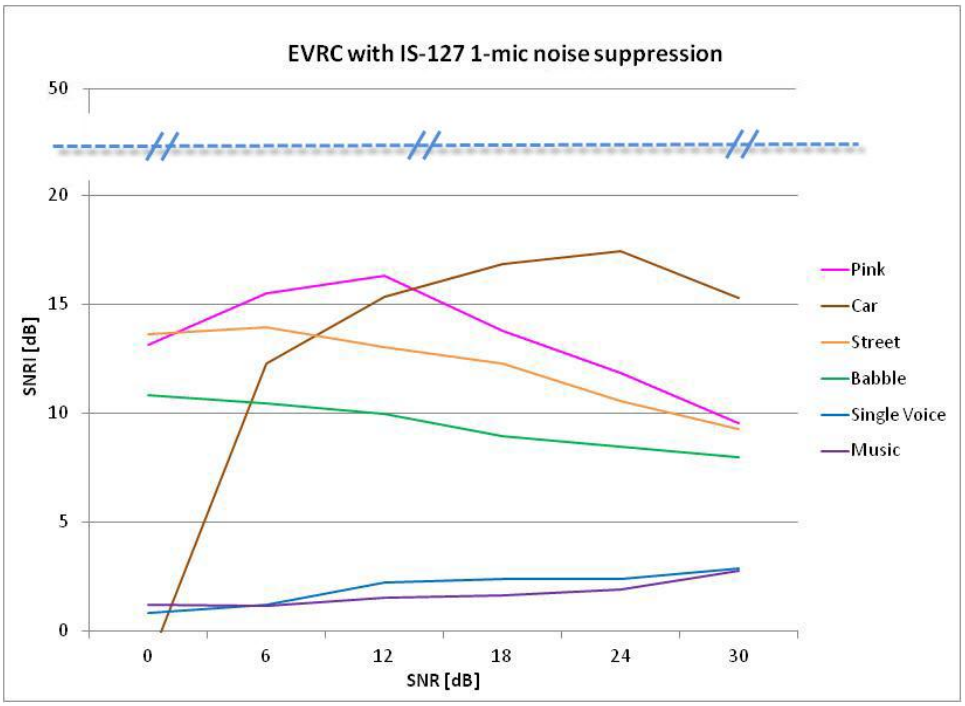

Figure 2. EVRC with IS-127 1-mic noise suppression

For the stationary noise types, pink, car, street, and to some extent babble, there is some noise reduction. However, there is essentially no noise reduction for the highly non-stationary noise types voice and music.

Figure 3 shows the results for the same conditions, but in this case the intrinsic noise suppressor in the EVRC is disabled and an advanced two-microphone noise suppressor is enabled. The noise reduction is greatly increased over the IS-127 suppressor at all levels of noise, and equally effective for the highly non-stationary noise types voice and music, which were not reduced at all by the IS-127 suppressor.

While not shown on these plots, it is important to note that good voice quality is preserved even at these levels of suppression.

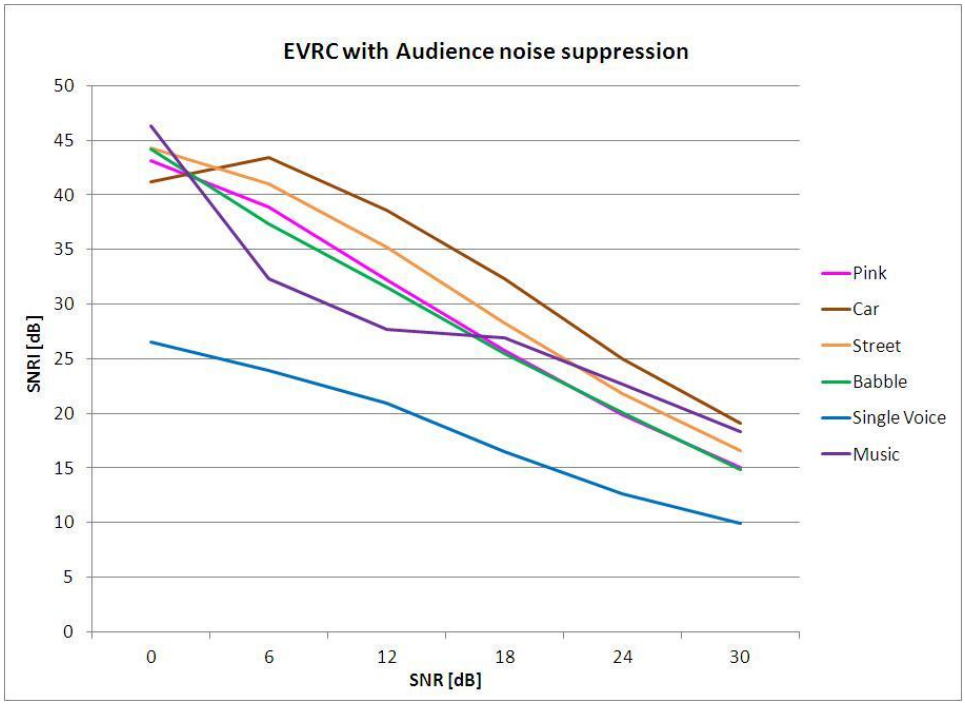

Figure 3. EVRC with Audience noise suppression

The reason so many device manufacturers and network operators are incorporating Audience earSmart technology into their devices and across their networks should be clear at this point. As we become a more mobile society, and as we make phone calls anywhere we want to, the issue of background and system noise will continue to hinder those on the other end of our calls from being able to hear every word of the conversation distinctly and clearly. With earSmart, the noise is reduced to a point of no longer being an annoyance, and clarity is restored to our calls.

There is another benefit for network operators as well. When a device is at the edge of a network cell, the signal from the handset back to the cell site is very weak and system noise can mean that the call is dropped or that when it is completed, it is difficult for the receiving party to hear. If the background noise of the conversation is filtered out effectively, then the signal is more intelligible when it is received by the cell site (and at a lower data rate) and routed through the network to the receiving party. If both the caller and the receiving party have devices with earSmart technology embedded into them, the end-

to-end conversation is better on both ends, fewer calls are dropped, and in some cases this will result in a savings to the network operator that might, otherwise, have to fill in some of the coverage areas with additional cell sites, an expensive and time-consuming process because the calls can be completed using a lower data rate.

Voice quality enhancements and noise suppression are becoming the accepted way of maximizing voice intelligibility and controlling background noise. As wideband audio is deployed in more networks, and as users are moving from a wideband-capable cell site back to a narrowband voice site (moving from 3G back to 2G for example), which happens frequently especially in the suburban and more rural areas of wireless coverage, the difference in voice quality will not be as noticeable to each of the parties on the call. Audience is the world leader in this technology and is involved in helping determine standards in noise suppression that will become ubiquitous in the wireless industry.

Network Capacity

An added benefit of noise suppression is the fact that within a given cell sector more calls can be handled because each call requires lower data rates therefore the capacity of the network is improved. This is an important step since there are very few ways in which to increase cell sector capacity. You can add more spectrum to a cell site, or build new sites closer together, but most network operators do not, at this point, have any more spectrum available to them, and building sites closer together is a time-consuming and expensive process. Therefore, being able to increase capacity by invoking noise suppression is of great interest to network operators.

Further, since many operators are beginning to convert their voice systems to wideband voice in order to increase the clarity of the calls, the number of calls that can be handled within a cell sector will actually decrease. By adding noise suppression, some of the capacity can be regained. It is expected that many more network operators and handset vendors will be embracing intelligent noise suppression technologies going forward for these two reasons.

Audience earSmart technology has been demonstrated to add capacity to a given cell sector on a network. As can be seen from Figure 4 below, the use of this technology reduces the amount of activity (therefore capacity). The baseline activity for clean speech, without noise suppression, is about 60% (that is, for the one minute of signal, about 60% contains active speech).

When the noise is present without noise suppression (solid lines), the activity increases, with the amount depending on how much and which type of noise is in the call. For only low amounts (30dB SNR) of stationary (pink, car, street) noise, the increase is small (roughly 5%). But for non-stationary sounds such as music, voice, or babble, even at the low levels (30dB SNR), the data rate increases to 80-87%. With noise suppression, the data rate is kept low, even slightly below the ‘quiet/no NS’ baseline. The next plot shows the reduction in activity due to NS.

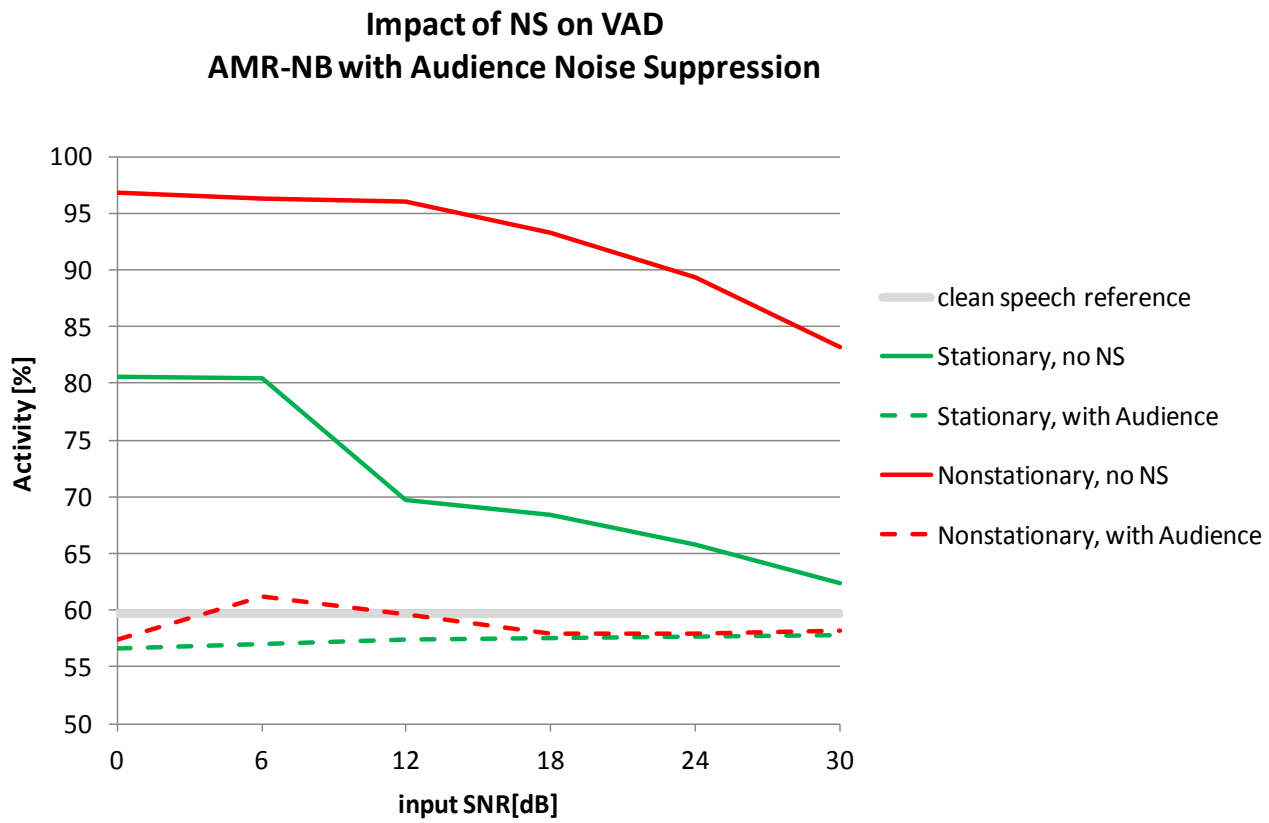

Figure 4. Impact of NS on VAD, AMR-NB with Audience Noise Suppression

This plot shows the VAD (Voice Activity Detection) improvement (or data rate savings) due to the presence of NS. From the previous plot, at 30dB SNR, the savings range from a low of about 5% for stationary noise to 20%-27%for non-stationary noise. Assuming that typical noise usage will fall in the range of 12 to 24 dB, the reduction in activity across noise types ranges from about 7%-12% (for pink, car, street) to 30%-36% (for babble, music, voice).

Thus the increased capacity would vary from 5% to the 20%-27% reported. This makes a significant difference in the number of simultaneous calls that can be handled within a given cell sector. The higher the noise level (the lower the SNR), the greater the improvement in voice quality and cell sector capacity.

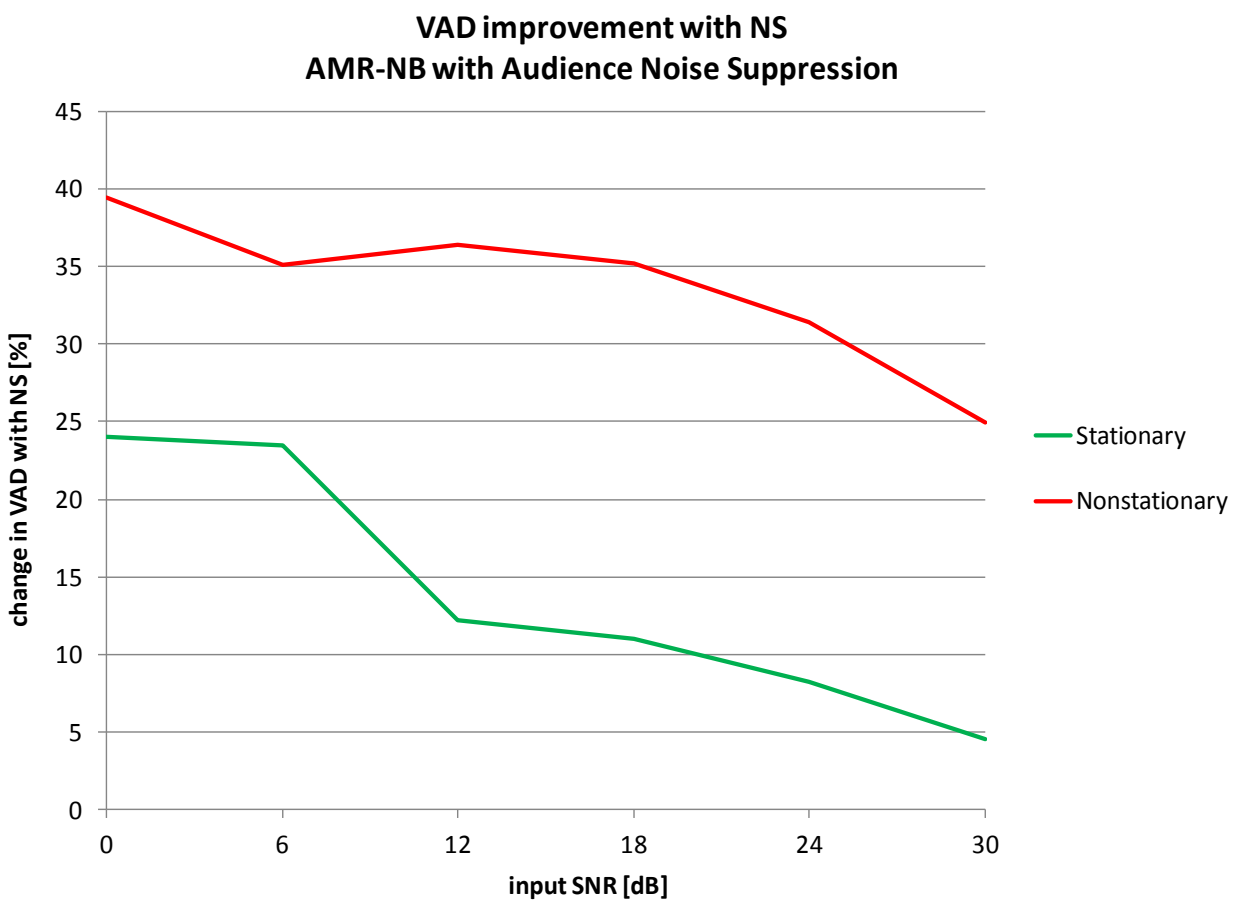

Figure 5. VAD improvements with NS, AMR-NB with Audience Noise Suppression

Averaging the savings for all noise types at 12, 18, 24 dB SNR gives a value of 22.4%.

Allowing for ~5% overhead, the average data rate savings can be roughly estimated at about 20%.

Assuming that 100% of the users in a sector are employing terminals equipped with Audience earSmart technology and are operating in a noise type and level in the range of the above, the working assumption is that that 20% reduction in data rate can be stated as a 20% savings in network capacity. This results in a huge savings to network operators in congested use areas. It is as though they gain 20% additional capacity simply by making use of earSmart technology in their devices.

There are multiple types of vocoders in use in the market so the figures below are similar plots for, AMR-WB (12.65 kbps), EVRC (SO3) and EVRC-B (SO68) at two Continuous Operating Points (COP=0 and COP=4) commonly used by CDMA network operators to manage loading. As can be seen, there is a significant improvement in noise level and thus the capacity required for each call. Therefore, Audience's technology is not only an advantage to the customer but also to the network operators. More capacity means more completed calls and more people within a cell sector who will have access to the network even in times of heavy congestion.

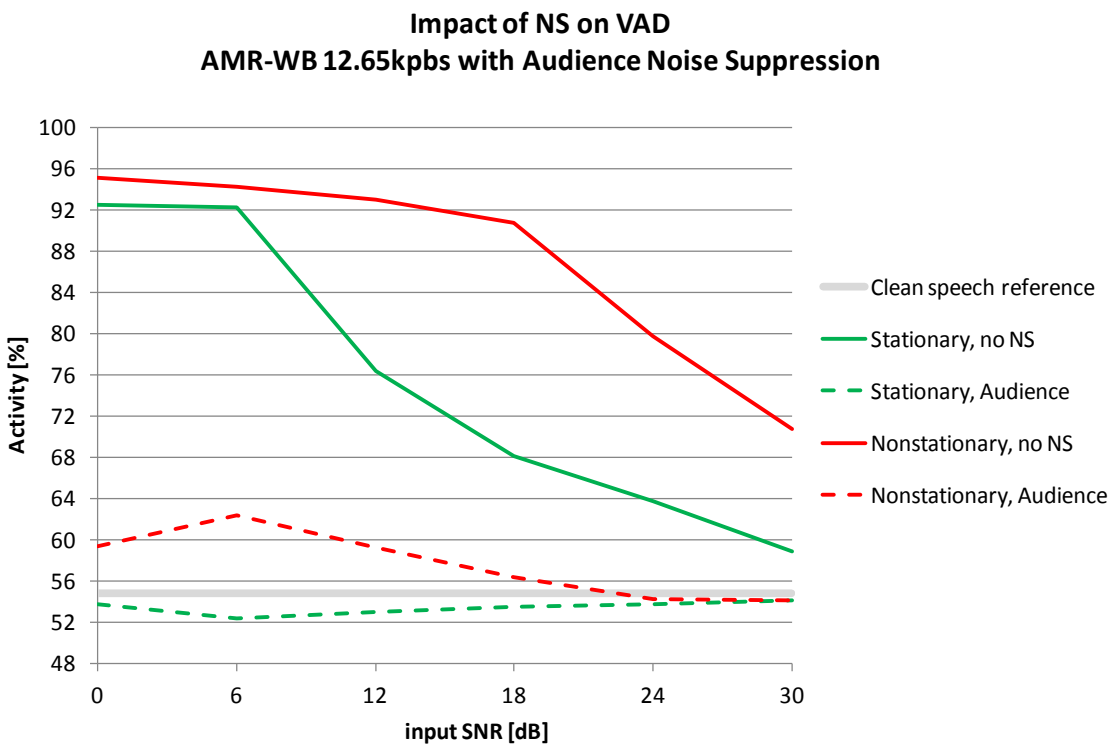

Figure 6 Impact of NS on Data Rate. AMR-WB

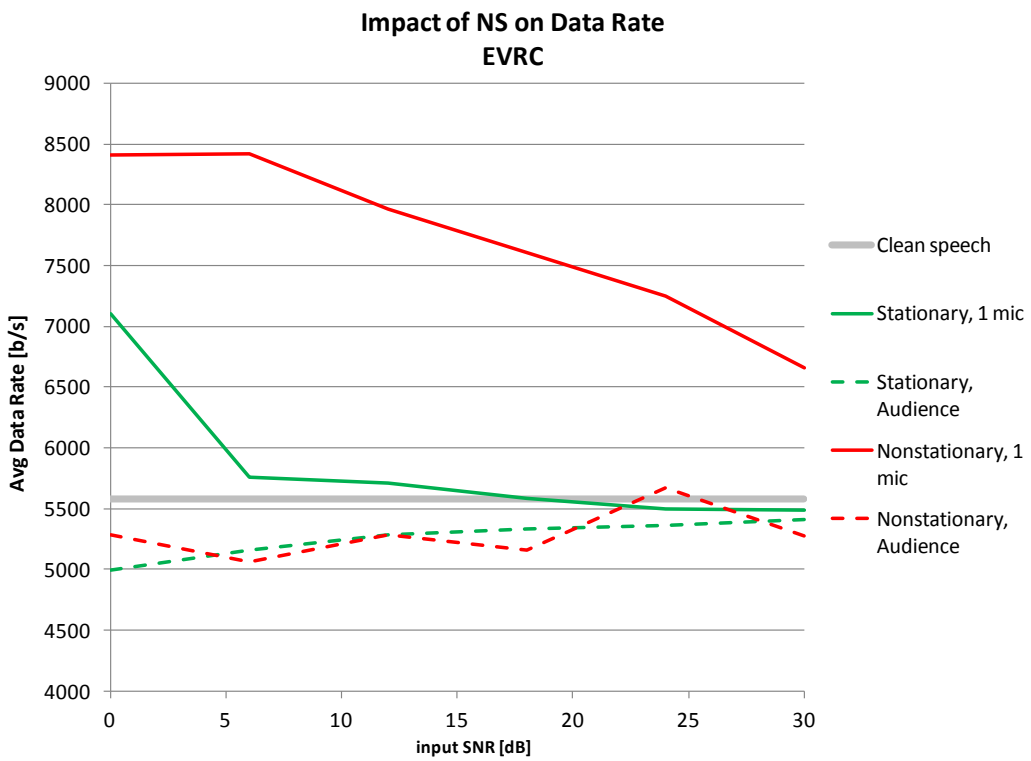

Figure 7. Impact of NS on Data Rate, EVRC

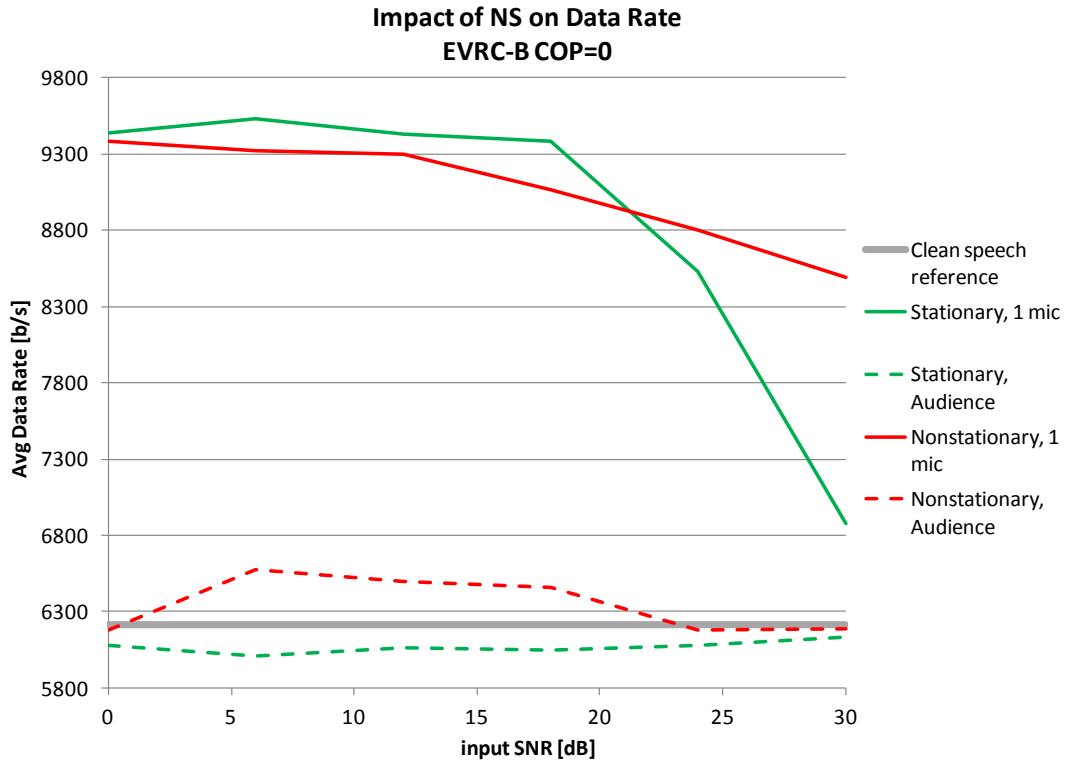

Figure 8. Impact of NS on Data Rate, EVRC-B COP=0

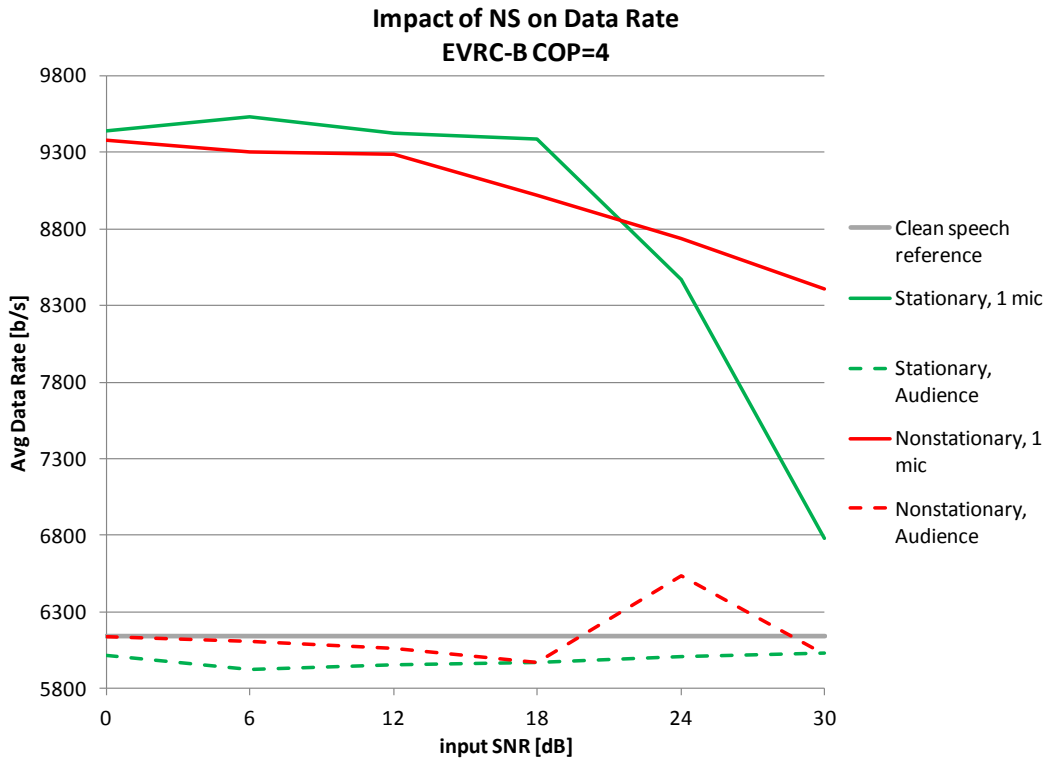

Figure 9. Impact of NS on Data Rate, EVRC-B COP=4

More Conversations per Cell Site

Another advantage of this type of interactive noise suppression within a broadband network is the fact that UMTS (WCDMA), HSPA, HSPA+, CMDA (EV-DO) and LTE make use of broadband as opposed to channelized spectrum. In the broadband world, each voice conversation is carried over the radio portion of the network using a specific code for each conversation and the receivers on each end see all of the other conversations as noise. As a phone moves from near the center of a cell toward the edge of the cell, the RF power of the mobile device is increased and the voice data rate is increased to compensate for the additional noise introduced into the receiver on the cell site end.

It is also important to note the trend from operators to transition networks from narrowband to wideband to provide improved voice quality. In wideband networks, voice is richer, as is noise, which makes noise suppression even more important. This network transition has been supported by the GSMA with the HD Voice specification, promoting advanced voice with noise suppression as the key feature.

If the voice portion of the transmission uses Audience technology, the voice quality is better and therefore more understandable even when the device is at the edge of the cell. The result is that customers have a better voice experience and and/or the network is capable of handling more voice calls per cell sector. This means that each cell sector can effectively handle more calls per sector, which in turn reduces the number of dropped calls and the number of calls that cannot be completed, which is a real benefit to network operators.

Conclusion

It is always exciting to find a company that identifies a problem or, in this case, multiple problems, assembles a staff of experts, and then finds ways to overcome the problems with smart products. Audience has taken a unique approach to solving several noise problems and has done so in an elegant and affordable manner. Many people have already experienced the solution Audience has created although many are not even aware that the technology is resident in their phones. They simply know that the phone works better in a noisy environment. Audience standardized a set of tests to gather background noise data, then set about solving the problems associated with minimizing the noise and providing a successful solution.

However, this is only the start. As we move into a world that will be dominated by wireless broadband for voice, data, and video, the technologies developed by Audience will advance to keep up with this new set of challenges. Meanwhile, what it has accomplished in reducing both background noise and noise generated with the radio spectrum has been well received. As Audience continues to move forward with its technology advancements we will all benefit. Network operators will have more tools at their disposal to provide their customers with a better wireless experience, and the wireless customers will enjoy the benefits of better voice quality and intelligibility. A true win for the industry.

Andrew M. Seybold