White Paper

The Importance of Mobile Audio
Sponsored by Qualcomm

SUMMARY

When it comes to smartphones, virtually everyone knows and cares about things like the number of megapixels for the onboard cameras, the resolution of the screen, what kind of video it can shoot, and just about everything else related to visual quality. Ask most people about the quality of the audio features, however, and you'll probably get little more than a quizzical look and shrug of the shoulders. That's unfortunate because in the same way that pairing a beautiful new 65" or larger 4K TV with a tiny, powered Bluetooth speaker as a sound system would ruin the overall media experience, not thinking about the differences that audio systems can make with smartphones doesn't make any sense either. If you want to enjoy the highest possible audio quality when listening to music, get the best possible audio response time while gaming and enjoy the highest quality voice calls while talking, then digging into the details of what a mobile sound system is capable of—including both the smartphone and any earbuds, headphones or speakers to which they are connected—is important to do. First, however, you need to understand a bit more about how digital audio and wireless audio works.

> "Audio quality is an incredibly important part of the overall AV experience that modern smartphones provide, but few people understand that the only way to maintain that quality is through an end-to-end focused technology solution."—Bob O'Donnell, Chief Analyst

INTRODUCTION

While some people view them primarily as mobile data terminals, today's smartphones are also tremendously powerful sources of audio-driven entertainment and communication. Streaming videos, listening to music, playing games and making phone calls, amongst many other activities, are incredibly important aspects of modern smartphone usage that are completely dependent on audio, extending from within the phone out to any connected listening devices. In the days when smartphones included headphone jacks, the "A" portion of the AV experience could be delivered through traditional wired headphones. As smartphone thickness shrank and people's desire for flexibility grew, however, the move to wireless audio delivered via Bluetooth became the de facto standard and that's where we stand today.

Bringing high quality wireless sound to life on mobile devices involves a lot more steps—and several more technologies—than most people realize, however, and many of these capabilities aren't widely understood. The goal of this white paper is to provide background and context on how digital audio works and offer a primer on what elements impact its quality.

Digital Audio Basics

Like digital images, digital audio is made up of large strings of 1s and 0s broken up into organized chunks of data in different sizes. Generally speaking, the more data you have, the more detail you have about the sound file it represents and, therefore, the higher the quality it contains.

Unlike digital images, sound is experienced over time, so the process of converting it from an audible, analog version of what we hear into a digital file is done through a time-slicing process called sampling. Essentially, from the moment the sampling process begins, a sound waveform is analyzed at tiny, regular intervals and the digital output of that process is recorded. The frequency at which that analysis occurs is referred to as the sampling rate and it's typically measured in thousands of times per second—common rates are 44.1 kHz (meaning 44,100 times per second), 48 kHz, 96 kHz, etc. The amount of digital data recorded at each of those sampling intervals is referred to as the sample size or bit depth with 16-bit and 24-bit being the most common.

The more samples that are taken per second and the larger the size of each sample, the more accurate the digital representation of the analog waveform will be. So, for example, 24-bit, 192 kHz audio has roughly 6.5x the amount of information as a 16-bit, 44.1 kHz audio sample (which happens to be the standard used for audio CDs or "CD quality") of the same waveform.

From a visual perspective, think of it like the difference between the image taken by an 8-megapixel camera versus one taken by a 50-megapixel camera.

As tempting as it may be to end the analogy there, it turns out human ears are not as sensitive as human eyes. Plus, because it is a temporal phenomenon, there's no such thing as zooming into an audio file. As a result, in some situations, having additional digital audio data resolution doesn't translate into a difference that people can actually hear. Instead, you meet a point of diminishing returns. The exact location of that point from a numeric perspective is the subject of raging debates among audiophiles, professional musicians and other interested parties (and, to be fair, varies from person to person because of physiological differences in hearing). The bottom line is that bigger numbers aren't always better from an actual experience perspective.

Discussions of digital audio quality need to cover much more than just basic sampling rate information, however. Digital sound and music files, like digital images, are also commonly subject to various forms of compression—things like MP3 format for music or JPEG for digital images. These compression techniques were created to reduce the size of the files involved, making them easier to store and distribute—a factor that was very important when these formats were created, but much less so now. The cost of storage has shrunk dramatically and today's WiFi 6 and 5G wireless networks can easily stream much larger files, so the economics of file transfer have evolved. As a result, it's possible to leverage new technologies that have been developed to take advantage of these changes.

In addition to file compression, real-time data like music and video are also subject to what I'll call stream compression when sent over a wireless link, in order to reduce the amount of bandwidth they require to be delivered. A piece of software called a codec (short for compressor-decompressor) reduces the number of bits necessary to stream from the mobile device to the headphones (or speaker) to ensure a reliable, consistent connection while still maintaining high audio quality.

Just to keep things interesting, you could have situations where you have music files that are recorded at high sampling rates, but then put into compressed filed formats before being streamed via one codec and others where the sampling rates are lower, but the files are uncompressed, and processed by a different codec and it would be extremely hard to tell what factors were ultimately influencing the quality of the signal (or the differences between them).

As if that wasn't enough to consider, the ability to hear differences in audio quality is also strongly dependent on the type of audio signal being listened to. For example, while you may be able to notice differences between, say, an uncompressed 24-bit, 96 kHz sample of a lone

singing voice and a compressed 16-bit, 48 kHz version, it would be much harder to do so with a heavily produced pop song.

The bottom line with mobile audio is it's best to start with and try to maintain the highest possible signal quality through the entire audio chain to avoid (or potentially offset) any quality impacts that occur along the way. In addition, it's critical to think about audio solutions from a complete end-to-end perspective.

COMPRESSION, CODECS AND WIRELESS TRANSMISSIONS

The most important factors in maintaining audio quality over a wireless link like Bluetooth are determined by the capabilities of the codec. The codec takes the audio streams, figures out methods to break the data into appropriately size chunks, applies compression technologies to the data, and then manages the process of sending (and potentially re-sending) these data packets to ensure a consistent signal is maintained. In an ideal environment, the process can be very straightforward, but in congested real-world environments like train stations, airports and other crowded places, interference from WiFi signals--many of which are transmitted at the same 2.4 Ghz frequency used by Bluetooth--and the varying amount of available bandwidth at a given moment can make the process extremely challenging.

There are also critical psycho-acoustic factors that have to be taken into consideration. It turns out, for example, that people are much more distracted by audio dropouts than a drop in audio quality for a brief period of time, so codecs need to be tuned to adjust the right parameters at the right time to maintain a continuous flow of sound. Psychoacoustic principles are also used to determine what elements of a given signal can be more compressed than others with the smallest audible impact in both file compression formats like MP3 as well as in wireless codecs.

One of the other big challenges for wireless smartphone audio is that the initial default codec chosen for Bluetooth—called SBC (low complexity sub-band codec)—has some fairly serious limitations. In particular, concerns around bandwidth, processing and power requirements, all of which were strongly influenced by limitations of the smartphones being produced nearly 15 years ago when some of these technologies were first introduced, limit the modern effectiveness of this original Bluetooth audio codec Unfortunately, usage of SBC continues to be very high and its widespread influence as the lowest common denominator across all Bluetooth devices has kept the bar for audio quality over Bluetooth connections lower than it should be.

Part of the problem is basic math. Data transfer rates for audio can be calculated by multiplying the sample rate of a given audio file by its bit depth (or sample size) by the number of channels. So, for example, streaming an uncompressed, CD-quality, stereo 16-bit,

44.1 kHz music file requires 1,411 kilobits per second (kbps)-44.1 kHz x 16 x 2 = 1,411. SBC, however, is limited to a maximum bit rate of 345 kbps—less than ¼ of what uncompressed, CD-quality audio requires—and many applications use half that rate or less. Even worse, SBC is unable to adjust data packet sizes based on the conditions encounters, such as multiple other Bluetooth users/devices, WiFi signal interference, etc. As a result, it ends up having to resend data packets multiple times over and that, in turn, translates into latency, or delays, in receiving an audio signal. For something like mobile gaming, those latency issues can become a huge problem and even for things like watching videos can cause annoying audio sync issues.

Newer codecs like Qualcomm's aptX Adaptive, for example, on the other hand both support higher transfer rates (up to 860 Kbps for stereo files) as well as the ability to adjust how packets are transferred based on existing conditions—both of which translate into noticeable improvements in real-world audio quality and reductions in audio latency. It's important to note, by the way, that even though the bitrate is lower than what's required for uncompressed, CD-quality audio, important improvements in the compression algorithms allow aptX Adaptive to handle up to 24-bit, 96 kHz audio files. Another benefit of the way aptX Adaptive works is that it can keep latency down to as low as 89 milliseconds, which is a significant improvement over 175-250 millisecond latencies often found with SBC codec-based devices.

WORKING AS A SYSTEM

While it's critical to have individual technologies that function well on their own, the final and arguably most important piece in enabling high-quality, wireless mobile audio is coordinating multiple elements to work together towards the goal. At the most basic level, this translates into having Bluetooth earbuds, headphones and/or speakers that also support the most advanced new codecs, such as aptX Adaptive. Because of how Bluetooth audio connections are designed to work, if a smartphone and a paired Bluetooth device can't find a shared, high-level codec during the pairing process, they will default back down to SBC as the lowest common denominator. This, in turn, will limit the potential quality of the audio connection. So, step one is to pair two devices that support the best possible shared codecs.

Beyond the simple pairing, however, there are system-level advances that also have to be considered. As discussed previously, one of the biggest challenges for maintaining high-quality, consistent and long-range Bluetooth connections is to avoid interference from 2.4 GHz WiFi signals. The best way to do that is to have the WiFi and Bluetooth subsystems and radios built into a mobile phone essentially coordinate with each other in order to reduce interference. If that isn't done, in "noisy" environments it can take as much as 15-20 retransmissions of an individual Bluetooth audio packet to transfer successfully, which can

lead to both latency and drop-out problems. If, on the other hand, a single vendor can provide solutions that were intentionally designed to co-exist across an integrated communications system, those problems can be avoided.

In addition, the ability to extend certain technologies, such as Qualcomm's Hi-Speed extension to Bluetooth, gives phones that integrate them the ability to achieve higher data rates and bigger link budgets (thereby increasing range and reliability) while still peacefully co-existing with a phone's WiFi system. Again, something like this is only possible with components from a single vendor that are designed to work together.

REAL-WORLD APPLICATIONS

The most obvious benefit of maintaining a high-quality wireless mobile audio system is producing higher quality sound on a more consistent basis. Practically speaking, this translates to things like being able to enjoy the full capabilities of high-resolution streaming audio services, such as Amazon Music's HD Ultra, as well as any pre-recorded high resolution audio files. For video streaming, this translates into higher quality, more impactful soundtracks, as well as reductions in potential audio latency or synchronization issues. For mobile gaming, the reductions in audio latency can literally mean the difference between life and death—for your character, that is—as human brains react more quickly to audio cues than they do visual ones.

In addition to these media applications, another critical but often overlooked aspect of smartphones is the quality of voice calls. Traditional Bluetooth connections limit your choice of codecs when making voice calls. Even improvements to voice services, sometimes called HD Voice, or wideband voice, that first started with VoLTE (Voice over LTE) can still only use the EVS (Enhanced Voice Services) codec. While that's a big improvement over old-school cell phone calls, it's still not as good as it could be. By leveraging a codec like aptX Adaptive, voice calls can be improved to what's called super wideband voice. In addition to improvements in audio quality, super wideband voice offers the ability to leverage noise cancellation technologies as well.

LOOKING AHEAD

Audio is clearly becoming an increasingly important part of how we use our smartphones, hence the need for the increased focus on the audio subsystems. It's also why there's been a lot of attention focused on the upcoming Bluetooth LE Audio standard, which is expected to be ratified later this year and supported in devices either later this year or early next year. In

the interim, however, there's still a great deal of work being done to improve the wireless audio experience with new smartphones.

In fact, some of capabilities promised for LE Audio—such as true wireless mirroring—are being delivered by individual vendors already, albeit with their own proprietary solutions. In addition, while LE Audio addresses things like improving baseline transfer rates for wireless audio, there are some elements for increased Bluetooth range and decreased interference that can only be addressed through system-wide coordination, which is clearly beyond the scope of the new LE Audio spec.

Ultimately, improving audio quality on smartphones is a critical step that virtually anyone can benefit from. Best of all, as great as these kinds of enhancements are now, there's still yet more to come. Building flexible codecs with enhanced data transfer rates also opens future possibilities for things like surround sound-style, multi-channel audio simulation, directional audio for things like AR and VR, and much more. The truth is, though audio may be seen by some as the "poor stepchild" compared to video when it comes to entertainment and content consumption, it has an incredibly important impact on how we experience and use our mobile devices. In fact, for many consumers, higher quality audio is one of the most requested features they have when it comes to their next smartphone purchase. That's why seeing advances and improvements in how wireless audio is delivered on mobile devices is so critical to us all.