DAT330 – Principles of Digital Audio
Cogswell Polytechnical College
Spring 2009

Week 1 – Class Notes

Sound, Numbers, and the Fundamentals of Digital Audio

1. Sound and Numbers

Physics of Sound

Acoustics - study of sound and is concerned with the generation, transmission, and reception of sound waves. The circumstances for the three phenomena are created when energy causes a disturbance in a medium. This disturbance makes an object vibrate which in turn will generate a back and forth motion of the air molecules surrounding the object. The disturbance creates regions of pressure above or below normal atmospheric pressure.

Nodes and Anti-nodes – region of maximum displacement and region of minimum displacement, respectively.

Sound propagation and transmission.

Transducers – converts energy from one form to another (microphones and loudspeakers).

Vibration (periodic/ aperiodic) and Frequency

Hertz (Hz) - number of vibration cycles that pass a given point each second.

Wave components and properties – frequency, period, amplitude, wavelength, phase, sound velocity, diffraction, refraction, and reflection.

Sound Pressure Level

Sound pressure displacement above and below the equilibrium atmospheric level is described by the amplitude and is measured with the sound pressure level (SPL). The unit is given by the decibel (dB). The decibel is defined as 10 times the logarithm of a power ratio:

$$\mathit{Intensity\ Level} = 10\log\left(\frac{P1}{P2}\right)^2 \ dB$$

$$\mathit{number\ of\ decibels} = 10\log\left(\frac{I1}{I2}\right) = 10\log\left(\frac{P1}{P2}\right)^2 = 20\log\left(\frac{P1}{P2}\right)$$

where *P1* and *P2* are values of acoustical or electrical power.

In acoustic measurements, an intensity level (IL) can be measured in dB by setting the reference intensity to the threshold of hearing (10^{-12} W/m²) such that $P2 = 10^{-12}$ W/m².

Threshold of Human Hearing – 0 - 120 SPL

Harmonics

Sinusoid or sinewave – simplest form of periodic motion. Only has fundamental frequency.

All other periodic waveforms are complex because they contain a fundamental frequency and a series of frequencies at multiples of the fundamental.

Aperiodic complex waveforms.

Overtones and Timbre

Human frequency range – 20 Hz to 20 kHz

Fourier Theorem – complex periodic waveforms are composed by a harmonic series of sinewaves, such that they can be synthesized by adding sinewaves. Likewise, a complex periodic waveform can be decomposed into its sinusoidal components.

Digital Basics

Acoustic and analog audio technology are concerned with the continuous values of mathematical functions, on the other hand, digital audio systems deal with discrete values meaning that waveforms can be represented as numbers. Digital systems utilize a base 2 or binary representation of numbers since it facilitates the arithmetic and logic operations.

Number systems

Base 16 or Hexadecimal	*Base 10 or Decimal*	*Base 8 or Octal*	*Base 2 or Binary*
0	0	0	0000
1	1	1	0001
2	2	2	0010
3	3	3	0011
4	4	4	0100
5	5	5	0101
6	6	6	0110
7	7	7	0111
8	8	10	1000
9	9	11	1001
A	10	12	1010
B	11	13	1011
C	12	14	1100
D	13	15	1101
E	14	16	1110
F	15	17	1111

The Binary Number System

Base 2 or binary number system are more efficient for digital computers and equipment since only two numerals are needed to satisfy the machine's principal electrical concern of a voltage being turned on or off. This system can be easily represented by using 0 and 1; these binary digits are called bits.

Converting decimals to binary numbers: 2^n, where n = 0,1,2,...
Counting from right to left we find:

$$\ldots, 2^7, 2^6, 2^5, 2^4, 2^3, 2^2, 2^1, 2^0$$

For example, converting 5 to binary would be like so:
There is no 5 in a power of 2 number, but we can add such that we only need a 1 and a 4. Thus, we need to say "yes" on the rightmost column so that we have a 1, "no" on the next number for a 0, and "yes" again for the third column. Our resulting binary in a 4-bit word would be equivalent to 0101.

2^4	2^3	2^1	2^0
0	1	0	1

Hexadecimal conversion involves very much the same process, except for numbers above 10 we have to use letters. Usually, hexadecimal numbers are made up of two 4-bit words or an 8-bit byte. For example, 0110 1110 = 6E = 6,14

Binary Codes

Individual binary bits or numbers can be ordered into words with specific connotations attached, such that both the symbolic and numeric information can be easily encoded. Binary numbers are restricted such that an n-bit number can only encode 2^n numbers.

Negative numbers can be problematic because the sign must be encoded, with bits, as well. Signed-magnitude representation solves this issue by assigning a 1 in the left-most position and a 0 for a positive number, while the rest designate the value of the number.

Binary words – grouping of multiple 4-bit words to create a decimal digit:

$$a_3\, a_2\, a_1\, a_0 \quad b_3\, b_2\, b_1\, b_0 \ldots \quad n_3\, n_2\, n_1\, n_0$$

Weighted Binary Codes

Each bit is assigned a decimal value, called a weight. Each number represented by the binary weighted binary code is calculated from the sum of the weighted digits. For example, weights $w_3\, w_2\, w_1\, w_0$ and bits $a_3\, a_2\, a_1\, a_0$ would represent the decimal number
$N = w_3 \text{ x } a_3 + w_2 \text{ x } a_2 + w_1 \text{ x } a_1 + w_0 \text{ x } a_0$

Unweighted Binary Codes

Excess-3 code adds a 3 (0011) to each codeword such that each codeword has at least one 1. While 2-out-of-5 code is defined so that exactly two out of the five bits are 1 for every valid word. This definition is used as a simple way to check for errors.
Gray code, or reflected code system, only has one digit value change when counting from one state to the next.

Two's Complement

Simple binary systems can present problems when the result is stored because the left-most bit is a carry from the addition process that can be lost if the bit system is not large enough.

Forming a one's and two's complement in the binary system is easily done. Because the radix is 2, each bit of the binary number is subtracted from 1. Thus, a one's complement is formed by replacing 1s by 0s ad vice versa. The two's complement is formed by adding 1 to that number and observing any carry operations.

A positive two's complement number added to its negative value will always equal zero.

When handling both positive and negative numbers, that represent a waveform for example, the most significant bit (MSB) is the sign bit. When it is 0, the number is positive, and when it is 1, the number is negative.

Boolean Algebra

Boolean algebra is the method used to combine and manipulate binary signals. It provides the basis for decision making, condition testing, and performing logical operations. Using Boolean algebra, all logical decisions are performed with the two binary digits 1 and 0, a set of operators, and a number of laws and theorems.

Boolean operators:

The operators OR, AND, and XOR combine two binary digits to produce a single-digit result. NOT complements a binary digit. NAND and NOR are derived from the other operators. The operators can be used singly or in a combinational logic to perform any possible logical situation.

NOT - $F=\bar{X}$ complements any set of digits. If 0, the complement is 1 and vice versa.

AND - $F=X \cdot Y$ if X and Y are both 1, then the result is one; otherwise the result is 0.

OR - $F=X+Y$ if either X or Y, or both are 1 then the result is 1; otherwise the result is 0.

XOR - $F=X \circ Y$ if X and Y are different, the result is 1; otherwise 0, if X and Y are the same.

NAND - $F=\overline{X \cdot Y}$ combination of AND and NOT.

NOR - $F=\overline{X+Y}$ combination of OR and NOT.

Boolean operators can be combined to form meaningful expressions and lead to greater insight of the condition, or its simplification. Logical expressions correspond directly to networks of logic gates, realizable in hardware or software.

Truth tables, or table of combinations, can be used to illustrate all the possible combinations contained in an expression. The output is expressed in terms of the input variables.

Complementation, commutation, association, and distribution laws commonly used hold true for Boolean algebra. The use of Boolean theorems and other reduction theorems can simplify complex logical expressions.

Analog Versus Digital

Differences:

1) continuous vs. discrete values and representation of signals.
2) Signal chain artifact addition: noise and distortion for analog and digital circuits.
3) Cost and immunity to external agents (temperature, age, etc) of both.

2. Fundamentals of Digital Audio

Discrete Time Sampling

Digital systems use discrete numbers in order to do their operations. A continuous signal, such that of an analog waveform, must then be digitized by means of time sampling and quantization. In this manner, the infinitely variable signal is now represented as amplitude values in time. Discrete time sampling is thus the defining essential mechanism in digital audio systems.

It must be noted that information does not get lost, in “between” the samples, during the digitalization process if the signal is properly conditioned. The input signal is conditioned by applying a lowpass filter in order to remove frequencies that are too high to be properly sampled. Also, the output will be continuous because of the interpolation function that is used to recreate the signal. In conclusion, a signal with a finite frequency response can be sampled without any loss of information and can be completely reconstructed from these samples.

The sampling Theorem

The Shannon-Nyquist sampling theorem states that a continuous band-limited signal can be replaced by a discrete sequence of samples without loss of any information and describes how these samples are used to reconstruct the original continuous signal. The sampling frequency must be at least twice the highest signal frequency. For audio signals with frequencies ranging from 0 – $S/2$ Hz can be represented by sampling frequency, S samples per second. In general, the sampling frequency must then be at least twice the bandwidth of the sampled signal.

The Nyquist Frequency

Audio signals in order to be sampled have to be lowpass filtered so that their bandwidth is within the Nyquist frequency of $S/2$. Usually, this filter is designed so that frequencies lying beyond the human frequency range are completely removed. Now the signal is ready to be sampled.

The sampling frequency is defined as the number of samples per second. The sampling rate, the reciprocal of the sampling frequency, is the time between each sample.

For example, a waveform containing high frequencies will require a higher sampling frequency. A digital system's sampling frequency will determine the high frequency limit of the system and thus, its audio bandwidth.

If for some reason a signal has not been filtered and has frequencies above the Nyquist frequency, the signal will cause aliasing distortion.

The output of the system must also have a lowpass filter in order to remove high frequencies created within the system and is also used in the reconstruction process.

While higher sampling frequencies allow for a larger bandwidth of the input signal, we must note that higher sampling frequencies do not improve the fidelity of those signals within the bandlimited range.

The lowpass filters employed during the sampling process do no have a sharp attenuation at the cutoff frequency. Instead, they have a guard band in which their cutoff frequency is located well below the Nyquist frequency in order to ensure proper frequency attenuation.

Different sampling frequencies can be used, from 8 kHz to 192 kHz, depending on the application. However, the sampling frequency employed will require particular digital circuitry speeds, storage capacity, and medium transmission. For example, larger sampling frequencies require faster operating circuits and large amounts of data processing.

Aliasing

Aliasing, or foldover, occurs when the input signal has frequency components above the Nyquist frequency. In this case, the signal is interpreted in a different way than than that of the original signal. An increase in audio frequency will result in a decrease of sample point per period. As a result, aliasing produces unwanted frequencies in our signal. As the audio frequency increases, the alias frequencies will decrease. The relationship is as follows:

$$F_f = \pm NS \pm F$$

where S is the sampling frequency, F is the frequency above the Nyquist frequency, N is an integer, and F_f is the new frequency.

In the frequency domain, alias frequencies are shown as frequencies that folded over from the spectral images of the signal.

Quantization

Quantization represents the values of the measured amplitude of a waveform at sample time. Quantization determines the resolution of the characterization. When an analog waveform is sampled into pulses; the amplitude of each pulse yields a number representing the analog value at that instant. The accuracy of this measurement is limited by the system's resolution. Finite word lengths limit quantization resolution and introduce measuring errors. This error can be considered akin to the noise floor of an analog system.

An analog signal that is uniformly quantized will be mapped to a finite number of quanta of equal size. Because the infinite number of amplitude points present in an analog waveform have to be quantized by a finite number of quanta levels, an error will be introduced.

High-quality representations require a large number of quantization levels or number of bits in the quantizing word.

Signal-to-error Ratio

The word length of a binary system will determine the number of quantizing intervals available. An n-bit word will yield 2^n quantization levels, meaning that the larger the number of bits the better the approximation. However, because of the finite number of amplitude levels in a binary word, an infinitely changing amplitude signal cannot be digitized without error.

If a quantized interval lies outside of the analog value, either above or below, the error will be in the last significant bit of the quantization word. In other words, the binary number can not accommodate fractional values and thus, has to round up or down to the next available value.

Quantization error is defined as the difference between the actual analog value at sample time and the selected quantization interval value. The amplitude value will be rounded off to the nearest interval. The worst case will be when the waveform lies exactly between two intervals, such that the error is limited to $\pm Q/2$, where Q is a quantization interval (1 LSB). If a signal has large amplitude, the distortion will be proportionally small and likely masked. On the other hand, a small signal will have a large distortion and might be audible.

The ratio of the maximum expressible signal amplitude to the maximum quantization error determines the signal-to-error (S/E) ratio of the system. It can be considered in some ways similar to signal-to-noise (S/N) ratio of analog systems. The S/E ratio can be expressed as:

$$\frac{S}{E}(dB) = 6.02\mathrm{n} + 1.76$$

where *n* is the number of bits of the system. Longer word lengths increase the data signal bandwidth required to convey the signal. However, the signal-to-quantization noise power ratio increases exponentially with data signal bandwidth. Quantization error will be perceived as white noise for high amplitude signals (it is “seen” as uncorrelated distortion components), and as distortion for lower amplitude signals.

Quantization Distortion

Analysis of the quantization error of low-amplitude signals show that the spectrum is function of the input signal. The error will not be noise like since it is correlated. The reconstructed output signal will contain the in-band components of the error. The quantization error being a function of the original signal, will then be classified as distortion.

The magnitude of the error is independent of the input signal amplitude, but is dependent on the size of the quantization interval. Greater number of intervals will produce lower distortion. However, the number of intervals used to quantize the signal is also important. For example, a maximum peak-to-peak signal will utilize all quantization intervals. However, if lower signal levels are quantized, fewer quantization levels are used resulting in fewer intervals employed. This will be equivalent to using lower bit quantization settings.

Quantization distortion may add components above the Nyquist frequency and cause aliasing, or add harmonics to the signal and “reconstruct” the signal into another waveform if the input signal is simple. Other effects of quantization error is granulation noise and beat tones.

Dither

Dither is a small amount of uncorrelated noise that is introduced to the audio signal. It is added prior to sampling and aids in linearizing the quantization process by shifting the audio signal with respect of the quantization levels. In this case, the signal is uncorrelated by avoiding any periodicity in the quantization patterns in consecutive waveforms. Thus, any quantization errors are decorrelated from the signal by randomizing its effects to the point of elimination. However, this process will add noise to the output of the audio signal.

While dither does not completely mask quantization errors, it allows the digital system to encode low amplitude signals (smaller than the LSB).

Types of Dither

Dither signals are differentiated by their probability density function (pdf). Meaning that a random signal with a continuum of possible values will have a probability of values falling within a given interval. This probability defines the area under a function. For example, a dither signal might have equal probability of falling anywhere over an interval, or it might be more likely to fall in the middle of the interval. The interval may be 1 or 2 LSB's wide.

Three dither signals are used in audio applications: Gaussian pdf, rectangular pdf, and triangular pdf.

Dither signals can have a white spectrum, but it can also be modified so that the dither is weighted and the noise floor is reduced. Triangular pdf is preferred for most applications, but the other pdf's add less overall noise to the signal.