Audio Measurement Handbook

Bob Metzler

Audio precision ®

AUDIO MEASUREMENT HANDBOOK

Bob Metzler
Audio Precision, Inc.

ABOUT THIS BOOK

The AUDIO MEASUREMENT HANDBOOK is intended as practical, hands-on assistance for workers in all phases of the audio field. Its treatment of topics does not involve mathematics beyond simple algebra. Although published by Audio Precision, the measurement techniques described generally are not specific to Audio Precision instruments, but apply to any audio testing equipment. The first section describes basic tools and techniques and includes guides to aid in the selection of appropriate audio test instruments. The second section discusses the common environments for audio testing. The third section applies these techniques to various commonly-used types of audio equipment and describes the typical ranges of performance to expect in that equipment. Section Four is a glossary of specific audio terminology and major specifications used in the audio measurement field.

ABOUT THE COVER

The classic audio test instruments on the cover are from the "museum" at Audio Precision. Most are from the personal collection of Wayne Jones; all are fully functional. The Model 200 oscillator was the first product developed in the legendary Palo Alto garage by Hewlett and Packard. The black crackle, wooden cabinet General Radio Model 1932A and the HP 330 provided a generation of broadcast engineers with their first ability to measure distortion. The Tektronix TM 500 instruments were designed by the engineers who later founded Audio Precision.

ACKNOWLEDGEMENTS

Contents

Section 4
Audio Measurement Glossary

Section 1

Tools and Techniques

Introduction and Basic Tools

Most audio equipment testing is done on a stimulus-response basis. A signal of known characteristics is fed to the input of the device under test (DUT) and the output of the DUT is measured. The performance of the DUT is determined by degradation of the output signal from the known input signal. Often, sets of measurements are made as the stimulus is swept or stepped across the frequency spectrum or across an amplitude range, and the desired performance information is determined by the relationship between the corresponding set of output measurements.

The most common stimulus for audio testing is a sine wave. The sine wave is unique since it is the only signal to have all its energy concentrated at a single point in the frequency spectrum. It is therefore relatively simple to analyze test results from single sine wave testing. Multiple sine waves (usually two) are used for intermodulation distortion testing, and large numbers of multiple sine waves are used for some new testing techniques. White noise, pink noise, square waves, and impulses may also be used as stimulus for certain types of audio testing. It is also possible to make certain measurements using program material such as music or voice as stimulus.

Most audio measurements are made with an instrument which is basically an ac voltmeter. Filters are usually incorporated into the voltmeter for many of the required measurements. Some advanced audio measurements are made with spectrum analyzers, which today are likely to be FFT (Fast Fourier Transform) analyzers rather than sweeping heterodyne analyzers or wave analyzers. Certain acoustical measurements may be made with devices usually known as RTAs (Real Time Analyzers), which consist of a large number of bandpass filters at staggered frequencies, each with its own voltmeter, in one package with a common set of input terminals. One key distinguishing characteristic of audio measurements is the wide dynamic range involved. Oscilloscopes are not

widely used for quantitative measurements in audio work, since only gross defects such as clipping are easily visible via waveform display (time domain).

Sine Wave Generation

Sine wave generation for audio test purposes may be accomplished by a wide range of techniques. The oldest, and still lowest-distortion methods use resistor-capacitor (RC) oscillator circuits such as the state variable, Wien bridge, twin-tee, or bridged-tee configuration. Function generators may also be used for audio signal generation. Synthesizer techniques of two general categories are available at audio frequencies: frequency synthesizers and direct digital synthesis. Frequency synthesizers use quartz-based oscillators, digital frequency dividers, and phase-locked-loops to generate signals. Direct digital synthesis generates a digitally sampled version of a signal by sending binary words (either from a table in memory, or generated "on the fly" from an algorithm) to a digital-to-analog converter at a regular sample rate. Each of these techniques has its own advantages and disadvantages, as discussed below.

RC Oscillators

RC oscillators can be made with almost arbitrarily low distortion (0.0001% or better is achievable), and have no inherent spurious output frequencies. Their amplitude is not absolutely constant with frequency and so must be leveled by a another circuit for most test instrument applications. They are also subject to significant amplitude bounce and thus a need for settling time when the R or C value is switched to change the output frequency. Frequency accuracy depends on the R and C value tolerances, with values of several percent typical and better than one percent achievable.

Function Generators

Function generators have intrinsically constant amplitude versus frequency and they settle in amplitude almost immediately when frequency is changed. However, the basic waveform produced by a function generator is a triangle, not a sine wave. Function generators use resistor-diode networks or other nonlinear techniques to shape the triangle waveform into an approximation to a sine wave, but the remaining distortion is still too high (1% to 2% typical) for many audio applications. Frequency accuracy depends on R and C values.

Frequency Synthesizers

Frequency synthesizers can produce extremely accurate frequencies, since the output frequency is directly dependent on a quartz crystal oscillator. However, highly accurate frequencies are not needed for the vast majority of audio measurements. Synthesizers have several signals at different frequencies in dif-

ferent circuits, and spurious output signals are a problem with this architecture. Waveform distortion is not necessarily low, rarely below 0.1% to 1%.

Direct Digital Synthesizers (Waveform Synthesizers)

Direct digital synthesizers, sometimes called waveform synthesizers, can also produce accurate output frequencies since the sample rate clock can be related to a quartz-based oscillator. However, spurious signals higher than one-half the clock frequency can appear in the output and cause problems. Waveform distortion of this technique within the band depends principally on the resolution (number of bits) and accuracy of the D/A converters used to produce an analog output. Thus, a direct digital synthesizer with 16-bit D/A conversion can produce distortion directly comparable to compact disc players, at around 0.003%. If the out-of-band components are included as distortion, the typical performance is approximately 0.01%.

Current Technology Trends in Signal Generation

Where purity of the generated sine wave signal is the most important factor, RC oscillators are the best choice. The state-variable circuit is considered the best topology to achieve low distortion and rapid settling time. The limitation of relatively modest frequency accuracy is usually not a serious disadvantage for most audio measurements. It is also possible to add phase lock circuitry to lock the frequency to a synthesizer when higher accuracy is required. When more complex waveforms are desired, direct digital synthesis becomes more economical and advantageous. As higher and higher performance D/A converters become available which can run at the sample rates required for full spectrum audio testing, the direct digital synthesis technique is likely to become more common in audio test instruments.

Complete Audio Signal Generation Instruments

An audio signal generator consists of much more than the basic oscillator circuit. Generally speaking, audio test requirements may require an amplitude range from as high as one watt (+30 dBm) for headroom testing on line-level devices to as low as a few microvolts to simulate microphone outputs. The required dynamic range is about 120 dB. Both balanced and unbalanced output configurations are required in order to properly drive both balanced and unbalanced input circuitry of the Device Under Test (DUT). The generator output may need to be isolated from ground by a transformer or other floating technique to avoid introducing noise into the DUT. A selection of several output impedances may be required in order to test DUTs designed to be driven from different source impedances. The net result of these several requirements is that an audio test generator normally consists of a basic oscillator circuit followed by a power amplifier, attenuator and variable gain controls, balancing

device, and selectable output impedances; see Figure 1 for a simplified block diagram of a generic complete audio test generator.

While only sine wave generation has been discussed up to this point, complete audio test generators often also have the capability of generating more complex signals such as square waves, pairs of sine waves for intermodulation distortion testing, white and pink noise, controlled bursts of sine wave, and perhaps a variety of digitally-generated signals. sine wave bursts are often used for transient response tests of loudspeakers. Sine bursts are also useful for tests of compressors and limiters, as described in the Companders section starting on page 99.

Amplitude Calibration

Since sine waves are the most basic waveform and the rms (root mean square) technique is normally the most useful way of stating signal amplitude, sine wave audio generator output amplitude calibration is almost invariably in terms of the rms value of the waveform. For complex waveforms, it is most common to calibrate the generator output amplitude in terms of the rms value of a sine wave which would have the same peak value. For example, assume a sine wave generator is set for an indicated output level of 1.000 V (rms). As is well known, the peak value of the sine wave is then 1.414 V (2.828 V peak-to-peak). If any other waveform is selected at the same indicated 1.000 V setting, that other waveform will also have a peak amplitude of 1.414 V (2.828 V peak-to-peak). The advantage of this calibration philosophy is that if the sine wave amplitude was adjusted to be below the clipping level of the DUT, the device will then not clip when any other waveform is selected. But, it must be remembered that the indicated 1.000 V value is not the rms value of any waveform other than sine wave. If a square wave is selected, its rms value will be greater than the indicated value since a square wave has a crest factor (peak-
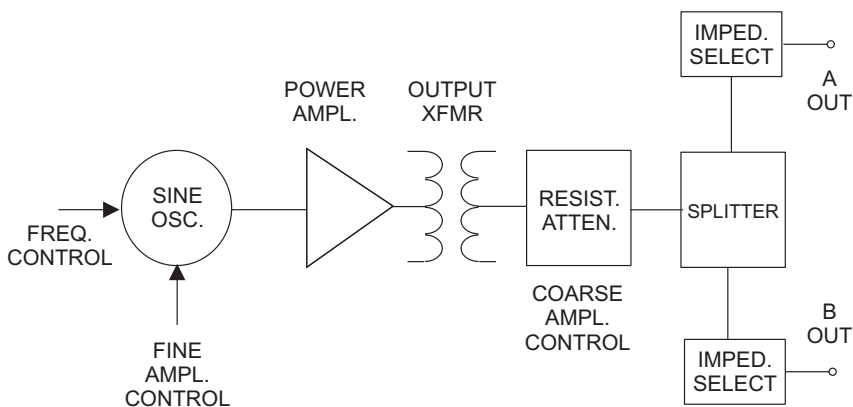
Figure 1. Block diagram, typical audio generator.

to-rms ratio) of 1.000. Most other complex waveforms will have rms values considerably less than 1.000, since the crest factor of most waveforms is greater than the 1.414 crest factor of a sine wave.

Important Generator Characteristics

The most important criteria for selecting among audio signal generators include frequency range, amplitude range, residual distortion, absolute accuracy of amplitude calibration, relative flatness of frequency response at any given amplitude setting, absolute accuracy of frequency calibration, and features such as selectable impedances, balanced/unbalanced configuration, floating-grounded configuration, etc. The better audio generators available in 1993 have frequency ranges extending from 10 or 20 Hz at the low end to 100 kHz or more at the high end, maximum amplitudes of +28 to +30 dBm, minimum amplitudes of –70 to –90 dBu, residual distortion from 0.0005% to 0.003%, absolute accuracy of 0.1 to 0.2 dB, and frequency response flatness of ±0.05 dB or better. The specific requirements will be determined by each specific application, but many engineers prefer to buy test equipment of the highest quality so that it can continue to serve their needs for many years as the state of the art improves.

Amplitude Measurements

Audio amplitude or level measurements are normally made with an ac voltmeter plus supporting circuitry. In its most basic form, an ac voltmeter consists of a detector and an indicator. The function of the detector is to convert the continuously-varying ac signal into a steady-state (dc) signal proportional to some parameter of the ac signal such as its rms (root-mean-square) value, its average value, or its peak value. The indicator then displays the dc value in a calibrated fashion. Older instruments used mechanical meters. More modern instruments typically use digital voltmeter technology, with the actual display being LED or LCD numeric readouts, or a computer display screen.

Rms Detectors

For most measurement applications, a true rms detector is preferred. Rms detection gives an output dc voltage proportional to the square root of the mean of the square of a series of instantaneous values of a varying waveform. The rms technique is the only detection technique whose output is proportional to the power (heating effect) of the signal, independent of the waveform of the signal. Thus, rms detectors are immune to phase shift effects and will provide accurate measurements for distorted sine wave signals, complex signals such as intermodulation test signals or voice and music, and noise signals. Most audio analyzers and audio voltmeters developed since approximately 1980 incorporate true rms detectors.

Note that generators of complex audio signals such as intermodulation test signals are normally calibrated in terms of the rms value of a sine wave with the same peak-to-peak value, as described in the generator section earlier. Therefore, the rms-measured value will not remain constant when measuring the output of a device while such a generator is switched from sine wave to a complex waveform. The table below shows the rms value of several common complex signals when they all have the same peak-to-peak value.

Waveform	P-P Value	rms Value
Sine	2.828	1.000
Square	2.828	1.414
2 tone 4:1	2.828	0.825
2 tone 1:1	2.828	0.707

Average-Responding Detectors

Prior to about 1980, true rms detection technology was difficult and expensive to implement and most audio meters used "rms calibrated average responding" detectors. Average detectors provide a dc output proportional to the average value of the rectified input signal. These detectors were calibrated to display the same value as a true rms detector when a sine wave signal was measured. With other waveforms, average-responding detectors will not produce the same indication as a true rms detector. Thus, average-responding detectors should not be used for accurate measurements of any signals other than sine waves. When measuring noise, average-responding detectors typically display a value one to two decibels lower than a true rms detector.

Peak Detectors

Peak detectors provide a dc output proportional to the peak value of the input ac signal. Use of a peak detector is appropriate when it is desired to measure signal amplitude with respect to potential clipping problems. Peak detectors are also useful when evaluating noise "spikes" of high amplitude but short duration. Such "spikes" may have little power due to their small area under the curve, but can be quite audible. Thus, peak detection may provide a measurement which correlates better to human audibility for this type of signal. A peak detector has an RC filtering circuit with a very short charge time constant and a very long discharge time constant. Practical peak detectors must have a finitely long charge time, even though the theoretical goal is zero. Similarly, they cannot practically have an infinite discharge time. Therefore, peak detectors do have some restrictions as to the repetition rate of input signal which they can accurately measure.

Detector Response Times

The minimum acceptable integration time for a detector is a function of the signal frequency and the desired accuracy. Lower frequencies require a longer integration time in order to make an accurate amplitude measurement. With numeric display instruments, including essentially all contemporary designs, the update rate of the display is dependent upon the detector integration time; there is no point in updating the display at periods much shorter than the detector integration time. Manually-operated instruments are normally designed with a single integration time long enough for the lowest frequency which the instrument is intended to measure, and a correspondingly-slow display update rate. These compromises are usually not serious when compared to human response times in noting a reading or taking any action based on the reading. In automatic test instruments, however, human response times need not be factored into most operations since the instrument and its controlling computer may manage many operations and ultimately present only final results. Therefore, sophisticated automatic audio analyzers may have a selection of integration times and display update rates available and may automatically select between them as a function of the frequency of the signal being measured. This automatic selection of response time is typically part of the overall algorithm described in the Measurement Dynamics and Reading Settling section on page 56.

Crest Factor

Crest factor is the ratio of the peak value to the rms value of a signal. A sine wave's crest factor is 1.414:1 and a non-bandwidth-limited square wave's crest factor is 1:1. More complex waveforms can have arbitrarily high crest factors, and true random noise has an infinite crest factor. Any detector has a maximum crest factor capability rating, above which it will not meet its stated accuracy. The crest factor rating of a detector assumes the peak input signal is at full scale. Reducing the signal amplitude below full scale improves the ability of the detector to handle high crest factors. In order to be useful in measuring complex signals and noise, an audio meter should have a crest factor rating of at least 5 at full scale; still-higher ratings are desirable. Note that the maximum crest factor rating of the detector portion of an instrument will not be realized if the signal is clipped in a preceding amplifier stage; see the discussion below in the Autoranging section.

Amplitude Measurement Units

The basic unit for electrical amplitude measurements is the volt (V). Volts, with millivolts (mV) and microvolts (μV) being the practical values for many typical audio levels, are commonly used units in certain parts of the audio field. Consumer high-fidelity and stereo equipment typically specifies device

input sensitivity and output levels of control amplifiers (where significant power is not being transferred) in volts.

Power

The basic unit for power measurements is the watt (W). Consumer and professional power amplifier output levels and loudspeaker capabilities are typically rated in watts. Measurement or computation of power requires knowledge of two parameters: voltage and current, or current and resistance, or voltage and resistance. In practice, audio power measurements are virtually never made by actual wattmeters. Normal practice is to measure the voltage across a known resistance value, and then to compute power from the V^2/R relationship. This computation may be automatic in computer-based instruments after operator entry of the known load resistance. Some instruments include power scales based on the instrument designer's assumption of a specific value of load resistance (usually 8 Ω), and will be in error for any other value of load.

Decibels

The majority of audio level measurements are made in some form of decibel unit. Decibels (dB) are a relative unit, and thus cannot communicate an absolute value until a reference value is stated. The common equations for decibel value computations from more basic units are:

$$dB = 20\log_{10}\frac{V}{V_r}$$

and

$$dB = 10\log_{10}\frac{P}{P_r}$$

where V and P indicate the actual measured voltage and power values and V_r and P_r indicate the reference values. Note that there are not two different kinds of decibels, voltage decibels and power decibels, as sometimes erroneously thought. The difference in the two equations is simply due to the fact that power is proportional to the square of voltage. For example, doubling the power in a given resistor (2:1 power ratio) requires increasing voltage by the square root of 2, or a 1.414:1 ratio. 10 log 2 is 3.01 dB, and 20 log 1.414 is also 3.01 dB.

Absolute Decibel Units

The reference value for decibels may sometimes be assumed rather than explicitly stated. For example, when a specification says "frequency response from 20 Hz to 20 kHz is ±2 dB" it means that the output amplitude at any frequency between 20 Hz and 20 kHz is within 2 dB of the output at 1 kHz; thus, the output at 1 kHz is the reference. Similarly, a statement that "signal-to-noise

ratio is 95 dB" means that the noise output of the device is 95 dB lower than the output under some reference signal condition.

If decibels are used to describe an absolute level, however, the reference value should be specifically defined. The most common absolute decibel units used in audio measurements are dBm, dBu, and dBV.

dBm

The dBm unit has been used for many years, principally in the broadcasting and professional audio sections of the audio field. The dBm units stands for decibels with reference to one milliwatt. Since watts are the unit of power, dBm is a power unit. Because audio meters are basically voltmeters, not watt-meters, the resistance across which the voltage is measured must be known before the dBm unit can be used meaningfully. The most common resistance (impedance) value in professional audio and broadcasting is 600 Ω, although 150 Ω is also common. Sometimes the dBm unit is written with the resistance reference as a subscript; for example, dBm $_{600}$. The measurement instrument must know or assume the resistance before it can display a measurement in dBm. Many simpler audio meters assume a value of 600 Ω as the circuit impedance. They then measure voltage and compute decibels from the ratio of the measured voltage to 0.7746 V, which is the voltage across a 600 Ω resistor when one milliwatt is being dissipated in that resistor. More sophisticated audio meters, especially if computer-based, may give the operator the ability to select among several common impedance references or even to freely enter any resistance value to be used as the reference.

Use of the dBm unit is improper if the circuit impedance value is not known, since the power level is then also unknown. Several decades ago, when professional audio technology was based on vacuum tubes and matching transformers, most items of audio equipment were operated on a maximum power transfer basis. This meant that the output impedance of each device was designed to be exactly equal to the input impedance of the following device, for matched connection and resulting maximum power transfer. Thus, a mixing console had a 600 Ω output impedance, the following line amplifier had a 600 Ω input impedance, and an audio meter bridged across the connection between them was correct in assuming power transfer with a 600 Ω reference. Modern solid-state audio equipment rarely operates on an impedance-matched power transfer basis. Output impedances are typically very low, below 50 Ω and sometimes near zero. Input impedances are relatively high, often 10 kΩ for professional equipment and 100 kΩ for consumer equipment. Thus, virtually no power is transferred, but the entire open-circuit voltage of the previous device is applied to the input of the following device.

However, old habits die hard. Furthermore, there is often a fiction maintained that 600 Ω is somehow significant. The driving device may be specified to work into a 600 Ω load, and the following device may be specified to be

driven from a 600 Ω source. But, there is usually no 600 Ω (or other known value) impedance present in either device or at the interface, and the dBm unit is thus usually not appropriate in modern audio equipment.

dBu

The dBu is a voltage-based decibel unit, referred to 0.7746 V. As noted above, 0.7746 V is the voltage across a 600 Ω resistor when dissipating one milliwatt. Thus, dBu and dBm are numerically equal in a 600 Ω circuit. But, the dBu unit assumes nothing about impedance. dBu is the proper unit to use in most professional audio and broadcasting applications unless it is definitely known that a matched condition exists, and the specific value of termination resistance is known. The audio voltmeters which have been manufactured for many decades are really dBu meters, even though the meter scale or panel switches may be labeled "dBm." These are voltmeters; they cannot measure current, resistance, or power, and therefore cannot possibly correctly indicate power. They are calibrated with the assumption of a 600 Ω circuit (0 dB = 0.7746 V) and should be labeled dBu meters, but often were produced before the dBu terminology became widespread.

dBV

The dBV is also a voltage-based unit, referred to 1.000 V. It is sometimes used in consumer audio equipment and less-often used in the professional audio and broadcast fields.
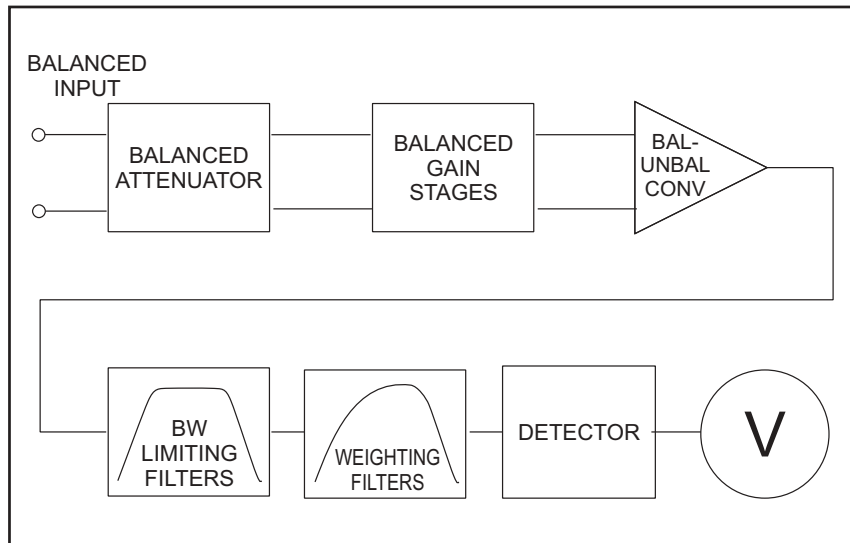
Figure 2. Audio level meter block diagram.

Complete Instruments

Just as a generator consists of more than an oscillator circuit, a complete audio meter consists of more than a detector and indicating device. For flexibility in a wide range of applications, the meter must have input circuitry compatible with measurement requirements. The instrument must be able to make accurate measurements over a wide dynamic range, typically ranging from several hundred volts maximum down to the microvolt level for noise measurements. A variety of filters are required for limiting noise bandwidth, rejecting interference, or weighting the effect of noise at different frequencies. Figure 2 is a generic block diagram of a flexible audio level meter.

The input impedance of an audio meter should be high enough to have a negligible loading effect on typical audio circuit source impedances. Since there are noise tradeoffs in creating extremely high input impedances, a value of 100 kΩ is a typical compromise for audio meter input impedance.

Termination and Load Resistors

For some applications, it is necessary to provide a specific termination resistance when testing a device. For example, a studio line amplifier may normally be connected to a 600 Ω load. During bench testing, it is desirable to simulate that same load, since a fundamental principal of test and measurement is to test devices under conditions as nearly equal to their actual usage conditions as possible. It is convenient if the audio meter has internal, selectable termination resistors for the most commonly-used values. This avoids the need for the user to locate an appropriate precision resistor of the correct value and to connect it in such a way that it does not introduce additional hum or noise into the measurement. For professional audio and broadcast applications, the value of 600 Ω is commonly specified. Germany and some of the Nordic countries may specify a 300 Ω load for some applications, and 150 Ω is also occasionally used. Power levels into these values of resistance will seldom be required to exceed +30 dBm (one watt), so internal termination resistors can be practical in an instrument.

When testing power amplifiers, however, it is not practical to provide termination resistors ("dummy loads") inside an instrument. Audio power amplifiers are commonly rated from tens of watts up to as much as thousands of watts. It is clearly not practical to provide for these levels of power dissipation inside a test instrument case. Thus, power amplifier testing is normally carried out with the amplifiers connected into external dummy loads which are typically heavy-duty wire-wound power resistors, often forced-air cooled. The audio meter input connects in parallel across the load resistor. If the meter is of the sophisticated variety which can compute and display power, there must be provisions for the user to enter the value of the external load resistors.
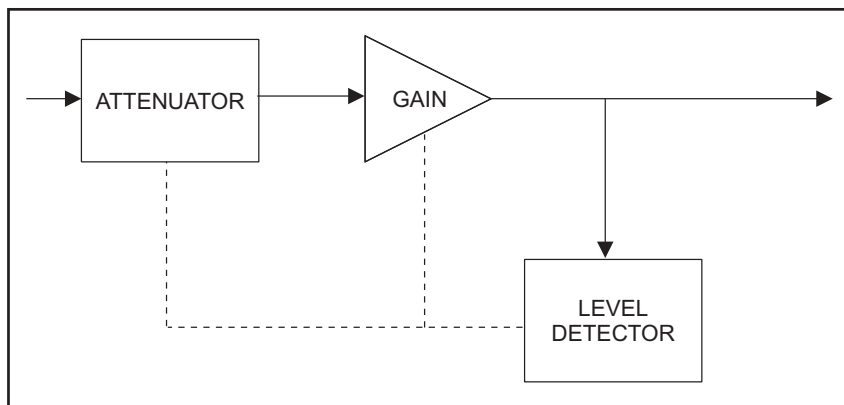
Figure 3. Block diagram, auto-ranging circuitry.

The input circuit of audio meters is normally capacitor-coupled so that any dc voltages present at the measurement point will not affect the measurements.

Balanced Inputs

High-quality audio meters normally have balanced inputs, consisting of two terminals both isolated from ground. The differential signal across those terminals is amplified and fed to the remainder of the instrument, while any common-mode signals present between either of the input terminals and ground is largely canceled out by a differential input stage. If the Device Under Test is connected to the meter by a cable with two conductors (twisted pair) under shield, any noise introduced into the cable tends to be in phase on both conductors and cancels out in the differential input stage. The ability of a balanced input stage to reject common-mode noise voltages while amplifying differential signals is called Common Mode Rejection Ratio (CMRR). CMRR values in excess of 60 dB across the 20 Hz–20 kHz audio spectrum are achievable with good design.

Ranging

To accept a wide range of input signals, audio meters normally have both switchable amplification and attenuation. On older and simpler meters, this range selection is accomplished manually by the operator selecting the range which gives the highest display without going off-scale. More modern and sophisticated meters are usually autoranging, where the optimum input amplification and/or attenuation is automatically selected so as to provide an output signal near the full-scale input of the detector. Over-driving the detector will cause measurement errors. Providing less than the optimum input level to the detector will cause the measurement to be more influenced by noise, and may also lead to longer settling times for the detector output to stabilize at a steady-state value.

Autoranging

Autoranging circuitry must have its own detector to measure signal level at the output of the amplification-attenuation ranging subsystem (see Figure 3). Most autoranging meters unfortunately use average-responding or rms-responding detectors to control the autoranging circuits. If the signal waveform has a high crest factor, such as white noise or a train of impulses, the average or rms level of the signal will be relatively low and an excessively sensitive input range will be automatically selected. This can result in clipping of the signal peaks in the amplifier stages before it gets to the measurement detector, causing measurement errors. These problems are even greater when filters are selected in the signal path. The only safe autoranging scheme uses a peak-responding detector to steer the ranging circuits, thus guaranteeing that there will be no clipping in the instrument.

No autoranging circuit can be completely effective with signals which have a very high peak-to-average or peak-to-rms value. Examples of such signals include impulsive noise ("spikes"), short bursts of signal repeated at very low repetition rates, single non-repeated events, and some types of music. In such cases, there is no automatic substitute for a knowledgeable user. The user must somehow determine or estimate what the highest peak signal will be, defeat the autoranging, and manually set the analyzer input range to a range which will not clip at the expected level. The sacrifice in increased error due to noise and possible slower operating speeds must be accepted as a compromise in such cases.

Accuracy and Resolution

All digital-display meters have accuracy and resolution specifications. The accuracy specification defines the possible error either as a percentage of instrument full scale or a percentage of the signal amplitude. The resolution specification is a further error term related to the finite "steps" of either the analog-to-digital conversion process or the display resolution itself. Measuring an unknown value and digitally displaying it with some fixed number of digits ultimately requires that the precise value of the signal be rounded off to the nearest value which can be displayed with the available number of digits (or converted with the available number of states of the converter). For example, a typical general-purpose digital dc voltmeter might have a specification of "0.5% ±1 digit" or "0.5% ±1 count" or "0.5% of reading ±0.05% full scale." The accuracy term (0.5% in these examples) is the dominant error term when measuring signals near full scale of the meter range in use. The resolution term (1 digit or 1 count or 0.05% full scale) will become the dominant error term as the signal amplitude approaches the bottom of the scale in use.

As an example, let's examine the worst-case specified error of a meter with a "3 1/2 digit display," meaning four digits but with the most significant digit only able to display a zero or a one. Thus, the significant figures 1999 are full

scale on any range. For simplicity of computations, let's assume a meter specification of 1% of reading ±1 digit. With a 1.5 V input signal on the "2 V" range (actually, 1.999 V full scale range), the 1% (of reading) accuracy term is 15 mV and the 1 digit term is 1 mV (one count in the least significant digit). At this level, the accuracy term heavily dominates the total error. With a 200 mV input signal, the 1% accuracy term computes to 2 mV and the resolution term is still 1 mV; at this level, accuracy and resolution are both contributing roughly equal amounts to the overall error. With a 20 mV signal still measured on the "2 V" range, the 1% accuracy term now becomes 0.2 mV, small compared to the 1 mV resolution term. We would say the meter is very resolution-limited under these latter conditions. In fact, if the signal amplitude were continuously reduced toward zero while fixed in the 2 V range, our ability to discriminate between values would be severely limited by the 1 mV resolution. A displayed 2 mV reading, given the ±1 count (1 mV on the 2 V scale) resolution, could actually be caused by a signal anywhere between 1 and 3 mV.

If the meter has another more-sensitive scale, such as "200 mV" (actually 199.9 mV full scale), the meter should be down-ranged as soon as the signal amplitude drops below 200 mV in order to minimize the resolution contribution to error. Ultimately, the meter will have a limiting "most-sensitive" range and signals which are a small fraction of full scale on the most sensitive range can be measured only with severe resolution limitations.

It should be noted that modern, automated, processor-based instruments often give the user no direct indication of which ranges they are using. For this reason, the trend is developing in such instruments to specify only an accuracy figure which includes all "percent of range" and offset effects.

The best available audio voltmeters have absolute accuracy specifications on the order of 1% to 2% (approximately 0.1 dB to 0.2 dB). Their most sensitive range should be such that resolution is better than 0.1 µV, so as to not contribute significantly to overall error when measuring noise levels approaching 1 µV.

Response Flatness

Absolute accuracy may be specified to include frequency effects. Even so, frequency response flatness of an audio meter may be important in addition to the accuracy specification, if the response flatness is better than the absolute accuracy. In many audio devices to be tested, flatness of frequency response may be the most important single parameter. If a meter manufacturer can specify response flatness of ±0.03 dB even though the best absolute accuracy which can be guaranteed is ±0.1 dB, this further specification can be useful to those engineers pushing the state of the art in frequency response flatness.

Noise Measurements

The audio meter discussion up until now has not included filtering of the input signal. For measurements of frequency response, power, level, gain, loss, and similar applications where the signal is a single sine wave with good signal-to-noise ratio (more than 40 dB), there is no need for filtering. In fact, filters represent an "operator trap" for such measurements.

For noise measurements (and distortion measurements, discussed below), a selection of filters is an absolute requirement. The word "noise" is sometimes used to refer to any undesired signal, including ac mains hum, stray magnetic fields from CRT monitors, etc. The following sections on noise measurement refer only to random noise where energy is distributed across a wide frequency spectrum, and not to undesired coherent signals such as hum and interference.

A measurement of random noise is not useful unless the measurement bandwidth is stated, due to the spectrally-spread nature of noise. It is not possible to accurately compare an actual noise measurement to a noise specification unless the noise meter uses the bandwidth (and possibly weighting filter) called out in the specification. For "white" noise (equal power per unit bandwidth), the measured power doubles (+3.01 dB) each time the measurement bandwidth is doubled.

Flexible noise-measurement instruments therefore normally include a selection of bandwidth-limiting filters. These may be available as actual bandpass filters, or as separately-selectable high-pass and low-pass filters. The most common bandwidth for noise measurements in professional audio, broadcasting, and consumer audio applications is 20 Hz–20 kHz, or the nearly-equivalent 22 Hz–22 kHz specified in CCIR 468. Communications applications are likely to specify a much narrower range consistent with the more limited bandwidth of communications-quality voice; 300 Hz–3.5 kHz is a frequent specification. Still other values of band-limiting filters may be offered for specific reasons. For example, some audio measurement instruments include a 400 Hz high-pass filter in order to provide much greater rejection of ac mains-related hum at 50 Hz or 60 Hz and the low-order harmonics of these frequencies. Some specifications refer to a "hum-to-hiss ratio," where the "hum" is a full 20 Hz–20 kHz bandwidth measurement and the "hiss" is a 400 Hz–20 kHz measurement which rejects the hum components. Other values of low-pass filters which are often provided include 30 kHz and 80 kHz; these selections are included in audio analyzers principally for distortion measurements rather than noise measurements, and will be discussed later.

Weighting (Psophometric) Filters

An entirely different category of filter used for noise measurements is the weighting filter or psophometric filter. The bandwidth-limiting filters described previously have flat frequency response over the center portion of the

AUDIO PRECISION   EQUAL LOUDNESS CONTOURS   LEVEL(dBr)   vs   FREQ(Hz)
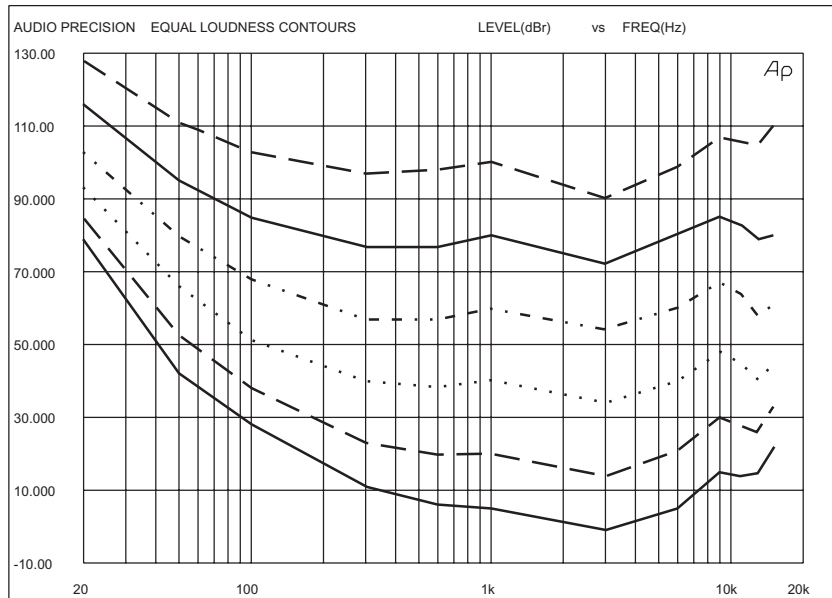
Figure 4. Fletcher-Munson curves.

audio band, between their upper and lower roll-off frequencies. Weighting fil-ters, on the other hand, have no flat response portion of their characteristics. Weighting filters are designed to provide measurements which correlate better with human perceptions than do unweighted measurements. It is well known that the sensitivity of the human ear is not flat with frequency, and furthermore that the frequency response of the human ear varies with the amplitude of the sound level. Figure 4 shows the classic Fletcher-Munson curves of typical hu-man ear sensitivity.

In order to obtain objective measurements which correlate well with human response, a number of different weighting filter curves have been proposed. Each is an approximation to the sensitivity curve of the ear under some condi-tions. Not all weighting filters are identical, for several reasons. Different sound pressure levels were assumed by different weighting filter developers, leading to use of different curves from the Fletcher-Munson family. One weighting filter (C message) includes the response curve of a typical telephone earpiece (of its day) in addition to the assumed human ear sensitivity curve. The CCIR weighting filter was deliberately designed to maximize its response to the types of impulsive noise often coupled into audio cables as they pass through telephone switching facilities. Virtually all weighting filters also repre-sent compromises between circuit simplicity and an ideal frequency response. Figure 5 shows response graphs of the ANSI A-weighting filter, CCIR 468 weighting filter, C-weighting filter, and the C message weighting filter.

A flexible audio meter designed for noise measurements will thus include a selection of weighting filters. Since various segments of the overall audio field

AUDIO PRECISION   FREQUENCY RESPONSE, FOUR COMMON NOISE WEIGHTING FILTERS

Figure 5. Weighting filter responses, actual measurements. Note that ANSI and C weighting filters are undefined above 20 kHz.

each tend to have their own weighting standards, it is common for meters to use a plug-in option strategy so that each user can buy and use only the filters needed in his particular specialty.

Time Variance

Noise, in addition to having energy spread over a broad frequency spectrum, is a time-variant signal. Thus, the measurement results will depend on the effective measurement time "window." Longer measurement times produce more integration of the noise signal, with the results thus being more repeatable from sample to sample. Shorter measurement times will produce more variable results, ultimately approaching instantaneous values. In addition to the integration intrinsic to long measurement periods, it is often desirable to do further statistical processing on the results of a number of consecutive measurement samples. It is not unusual to mathematically determine the maximum, average, and several standard deviation values (two sigma, three sigma, etc.) of a succession of noise measurements.

DUT Input Conditions for Noise Measurements

In many types of audio devices, most of the noise measured at the output is generated in the device input stage and amplified by the following stages. The level of noise generated in an amplifier is typically a function of the source im-

pedance of the amplifier. Therefore, noise measurements should specify an input termination condition, often called "back termination." The two most common specifications are with shorted input (0 Ω termination) and with a specific value of input termination, usually an approximation to the output impedance of the previous device normally connected to that point. For high-gain devices, the physical construction and shielding of the input termination used during noise measurements are very important to avoid erroneous measurements due to external signals coupling into the termination. Most modern high-quality audio generators are arranged so that their output connector is back-terminated with a resistor equal to the normal output impedance of the generator whenever the generator output signal is turned off. When testing very high-gain, low-noise devices such as microphone preamplifiers, however, it may still be necessary to disconnect the cable from generator to DUT input and replace it with a small well-shielded back termination directly at the DUT input, to avoid ground loops or the noise coupled into even well-shielded cables of any significant length.

Signal-to-Noise Ratio

Signal-to-Noise Ratio, often abbreviated as S/N or SNR, is simply a computation of the results of two measurements. The "signal" measurement is typically of a mid-band (usually 1 kHz) sine wave at a specified output level, usually the maximum rated or normal operating level of the device. The "noise" measurement must specify either the measurement bandwidth or a weighting filter. The ratio of the two measurements is then the S/N ratio of the device. This is normally considered to be the effective dynamic range of most types of device (but not for digital devices; see the digital audio measurements section on page 125 for a discussion of noise measurement in digital systems). S/N ratio measurements are greatly facilitated if the instrument features include a "relative dB" unit whose 0 dB reference can be set to equal the present value of input. With this feature, S/N ratio measurement becomes a simple sequence:

- Establish the specified output reference level.

- Press the button which causes this level to become the 0 dB reference.

- Remove the signal source and properly terminate the DUT input; may generators do this automatically at a button push. The meter indication is now the S/N ratio, although often expressed as a negative value (i.e., –90 dB for a S/N ratio of 90 dB).

Important Meter Characteristics for Noise Measurement

A number of the instrument characteristics already discussed are particularly important for noise measurements. The instrument must have a true rms detector available for most noise measurements. The CCIR 468 specification

requires another detector response, referred to as quasi-peak. The Dolby CCIR-ARM specification calls for an average-responding detector. The instrument should have a high crest factor capability, since noise can theoretically have an infinite crest factor. A selection of bandwidth-limiting and weighting filters should be available. The dynamic range and residual noise of the meter itself should be compatible with measurement of the lowest noise levels necessary for the particular application.

Signal Measurements in the Presence of Noise

The preceding sections have dealt with measuring noise, usually noise generated in the device under test. Conversely, it is sometimes desired to exclude noise from a measurement in order to accurately measure a signal near, or even below the wideband noise level. When the signal is a sine wave, a narrow tunable bandpass filter can accomplish this task. With the filter tuned to the signal frequency, the sine wave is passed without attenuation but noise outside the filter passband is attenuated. The narrower the filter bandwidth, the more broadband noise will be attenuated. However, as filter bandwidth is made narrower, it is more and more difficult to maintain accurate and stable gain through a filter capable of being tuned across a frequency range. Thus, there is a practical compromise between the error in measurement of sine wave amplitude caused by variable filter gain versus the potential noise rejection by bandwidth reduction. Tunable constant-Q (constant percentage bandwidth) filters of one-third octave bandwidth can be designed whose gain varies only a few tenths of a


Figure 6. Four pole one-third octave filter response.

decibel across a decade tuning range. A one-third octave filter has a –3dB bandwidth of about 20% of center frequency; see Figure 6 for a measured response curve of a four-pole one-third octave bandpass filter. If the noise to be excluded is approximately white noise, each 10:1 bandwidth reduction will result in a 10 dB noise power reduction. Thus, the 200 Hz filter bandwidth of a one-third octave filter tuned to 1 kHz will reduce measured white noise by about 20 dB compared to the standard 20 kHz audio bandwidth. This will permit reasonably accurate measurements of sine wave amplitude even when the sine wave level is as much as 10 dB below the 20 kHz bandwidth noise level.

Dynamics of Selective Measurements

When a selective bandpass filter is used to allow more accurate measurements of low-amplitude sine waves in the presence of noise, measurement dynamics are affected. A selective filter is an energy storage device, with the storage time inversely proportional to the filter bandwidth. A change in input level or tuned frequency produces a transient which takes time to stabilize. A one-third octave bandpass filter tuned to 20 Hz will take several hundred milliseconds before the filter output has stabilized to within 0.1 dB of the final value. Operators of manual test instruments normally wait until the instrument reading is stabilized before noting the value. Automatic test systems must automatically provide for that same result; see the section on Measurement Dynamics and Reading Settling on page 56.

Time Domain and Frequency Domain

Much of the remainder of this audio test and measurement concept discussion will depend upon an understanding of both time domain and frequency domain representations of signals. A time domain representation is a graph of amplitude versus time. This is the typical oscilloscope representation of a signal. A frequency domain representation is a graph of amplitude versus frequency. This is a spectrum analyzer representation of a signal. Any signal can be represented in both time and frequency domains. Certain phenomena are easier to understand in the time domain and others in the frequency domain.

A sine wave is familiar to most electronics technical people in the time domain (see Figure 7). The unique property of a sine wave is that all its energy occurs at a single frequency. Thus, the frequency domain representation of a sine wave is a single vertical line at the fundamental frequency (see Figure 8). Note that in this actual measurement, some harmonic energy is visible at third, fifth, and several higher-order harmonic frequencies due to non-linearity of the A/D converters used to acquire the sine wave.

A square wave is also a familiar waveform in the time domain (Figure 9). The square wave can be shown mathematically to consist of a fundamental sine wave and all odd harmonics, in phase, with the amplitudes of the harmon-

Figure 7. Time domain representation, 1 kHz sine wave.

Figure 8. Frequency domain representation, 1 kHz sine wave.

Figure 9. Time domain representation, 1 kHz square wave.

Figure 10. Frequency domain representation, 1 kHz square wave.

ics decreasing proportionally to the harmonic order. Thus, the third harmonic has an amplitude 1/3 that of the fundamental; the fifth harmonic is 1/5 the fundamental amplitude, etc. This is shown in the frequency domain representation of a square wave in Figure 10.

In general, any repetitive signal can be analyzed as a fundamental sine wave and harmonics of that fundamental frequency, with each harmonic having a specific amplitude and phase.

Measurement of Non-Linearity

A device under test with a completely linear transfer function (input-output relationship) will produce an output signal waveshape identical to the input waveshape, though possibly scaled up or down in amplitude according to the gain or loss of the DUT. If the transfer function is non-linear, the output waveshape will deviate from the input waveshape. In the frequency domain with a single sine wave stimulus, this non-linearity will show as energy at harmonics of the fundamental sine wave in addition to the fundamental itself (see Figure 11). Figure 8 above, as an example, shows third and fifth harmonic energy at 90 dB to 100 dB below the fundamental. If the test signal consists of two sine waves, a frequency domain analysis of the output will show both sine wave fundamentals, energy at the harmonics of both fundamentals, and also intermodulation distortion products at the sum and difference frequencies of the two original sine waves and of their harmonics (Figure 12).

Figure 11. Harmonic distortion concept.
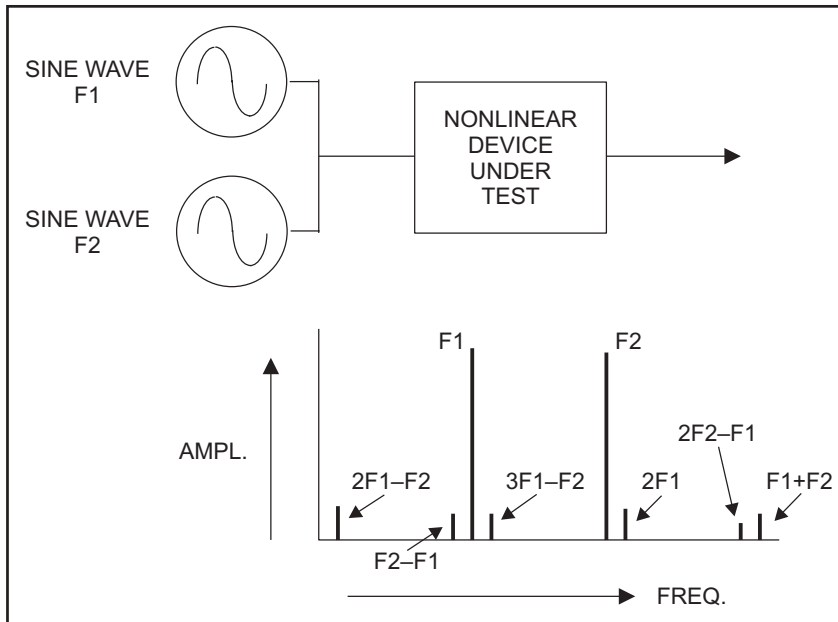
Figure 12. Intermodulation distortion concept.

Non-linearity can be measured by a number of different techniques. These include graphs of output amplitude versus input amplitude, measurements of total harmonic distortion plus noise (THD+N), total harmonic distortion excluding noise (THD), individual harmonic amplitudes, and a wide variety of different methods of measuring intermodulation distortion (IMD). The THD+N technique is by far the most widely used, and the vast majority of commercial distortion analyzers measure according to this technique.

Output Versus Input Amplitude

One straightforward method of measuring DUT linearity is by varying a sine wave stimulus signal across a wide amplitude range while measuring the output amplitude of the DUT. Ideally, the input stimulus will be varied from a maximum value at or above the overload point of the DUT down to a minimum value which produces DUT output approximately equal to the noise level of the DUT. Figure 13 shows a graph of output amplitude plotted vertically versus input amplitude plotted horizontally. On a nominally linear device, the graph over most of the range will be a straight diagonal line with one dB output increase per one dB input increase. At the top of the graph, it will become asymptotic to a horizontal straight line as the finite DUT output capability is reached ("clipping point") and the output can no longer follow increasing input levels. At the bottom of the graph, the output noise level of the DUT becomes a limitation and the graph will again approach a straight horizontal line at the noise level.

Figure 13. Input-output linearity graph.

The output-vs-input linearity measurement capability of an audio test instrument can be increased in the downwards direction by use of frequency selectivity in the analyzer. A narrow bandpass filter tuned to the sine wave stimulus frequency will reject most of the wideband noise generated in the DUT while still passing the sine wave without attenuation. Figure 14 shows the block diagram with measurement selectivity; Figure 15 shows the linearity curve of the same device as Figure 13, with the curve now extended downwards by the noise rejection capability of the filter. When the noise is approximately white
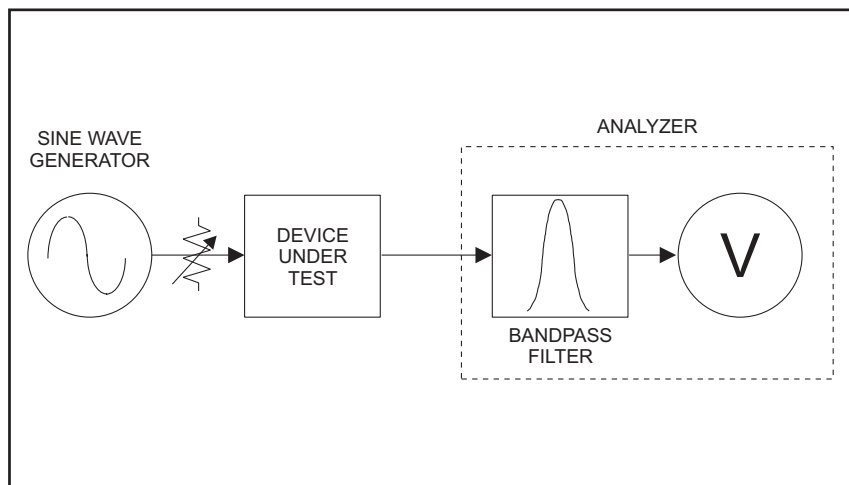
Figure 14. Block diagram, linearity measurement with selectivity.

AUDIO PRECISION   INPUT-OUTPUT LINEARITY, SELECTIVE MEASUREMENT

Figure 15. Input-output linearity with selective measurement.

in distribution and the available bandpass filter is tunable and of the constant-Q variety (constant percent bandwidth), it is advantageous to make the input-output linearity measurement at relatively low frequencies. The constant percentage bandwidth represents a smaller bandwidth in hertz at lower center frequencies and therefore accepts less noise. For example, a 1/3 octave filter tuned to 1 kHz is about 200 Hz wide; at 200 Hz, the same 1/3 octave filter is about 40 Hz wide. This 5:1 reduction in bandwidth will produce approximately a 7 dB reduction in noise if the noise is distributed in a white noise fashion across the spectrum. However, it may not be desirable to move the test frequency arbitrarily low for two reasons. First, stabilization time increases as the filter bandwidth becomes sharper and the detector integration time must also become longer at lower frequencies, so the overall measurement speed decreases. Second, many devices have some amount of residual ac mains-related hum at the ac mains frequency and second and third harmonics of the mains frequency. This hum is a coherent signal and will not be rejected if it falls within the bandpass of the analyzer filter. So, it is generally desirable to stay well away from the 50 Hz, 100 Hz, and 150 Hz area in most of the world and away from 60 Hz, 120 Hz, and 180 Hz in North America when making low-level measurements.

A graph spanning a 100 dB or 120 dB range as shown in Figures 13 and 15 is difficult to read with high resolution. Deviations from linearity of a few tenths of a decibel or even 1 or 2 dB cannot be read accurately. Some mathematical post-processing following the measurement can greatly improve the ability to detect small amounts of non-linearity. A practical technique involves

selecting some central portion of the measurement curve which is not near to either the clipping area at the top or the noise-limited area at the bottom. Line-fitting techniques such as the least-squares method can then be used to derive the straight line which best fits those measured values in the central region. If every measured point is then subtracted from this best straight line, the result is a curve of deviation from perfect linearity. This processed data can be plotted with a much more sensitive vertical scale, typically with zero at its center so that both positive and negative deviations from linearity can be plotted. Figure 16 shows the data from Figure 15 following this processing; now a vertical scale of ±5 dB allows small variations from perfect linearity to be easily seen.
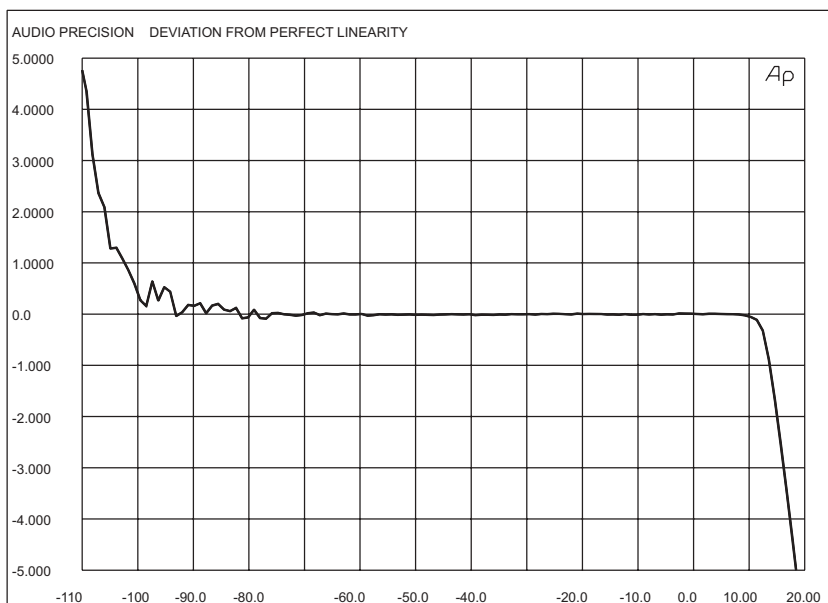
Figure 16. Deviation from perfect linearity.

   The output-vs-input linearity measurement method is very useful for examining the steady-state (static) characteristics of audio processing devices such as compressors and limiters, as will be discussed in Section 3. Figure 69 on page 100 shows the output-input curve of a compressor, where the compressor can be seen to be linear up to input values (horizontal axis) of about –37 dB, above which compression takes place. This technique is also very useful for testing D/A and A/D converters, where non-linearity at low amplitudes may be a significant performance shortcoming. This type of test will be further discussed in the digital converter section beginning on page 127.

Harmonic Distortion

Harmonic distortion can be measured in several different ways. All have the following basic concepts in common:

- Stimulate the DUT with a pure sine wave, so that for practical purposes there is no harmonic energy in the stimulus signal.

- Non-linearity of the DUT will cause harmonics of various order (2nd, 3rd, etc.) and various amplitudes to appear.

- Measure the output signal in such a way as to separate the harmonic energy from the fundamental energy. Typically, the harmonic energy is expressed relative to either the amplitude of the fundamental component at the output or to the amplitude of the complex signal (fundamental plus harmonics).

Individual Harmonic Distortion

A highly-selective tunable bandpass filter is required in order to measure individual harmonic distortion; see Figure 17. Some analyzers (often called wave analyzers or sweeping spectrum analyzers) achieve this effect with a superheterodyne receiver architecture, using a mixer, local oscillator, and fixed highly-selective filters at an intermediate frequency. Analyzers may also use a tunable filter directly sweepable through audio frequencies with no frequency conversions involved, or FFT (Fast Fourier Transform) analysis techniques to provide high-resolution analysis.

Figure 17. Individual harmonic distortion example.

Being able to isolate individual harmonics can be particularly useful to a design and development engineer, since different non-linearity mechanisms may give rise to different types of distortion. For example, transfer function non-linearities which are symmetrical around zero (such as the B-H curve of magnetic tape when recorded at high amplitudes; see Figure 19) produce odd harmonics. Non-linearities which are not symmetrical around zero (such as a section of a semiconductor junction transfer function; see Figure 18) produce dominantly even harmonics. Symmetrical clipping produces primarily odd harmonics, often of very high order. Thus, measuring the distribution of harmonics may provide good clues to where the distortion is being introduced.

Figure 18. Nonsymmetrical nonlinearity, resulting in even harmonics.

The performance limits of individual harmonic distortion measurement are typically poorer than other methods to be discussed below. To accurately measure a second harmonic component 90 dB below the fundamental, for example, calls for a bandpass filter with more than 90 dB rejection at one-half its center frequency; see Figure 20. Heterodyne-type sweeping spectrum analyzers can provide this selectivity, but their mixer stage limits dynamic range. FFT analyzers must be of at least 16-bit design to approach 90 dB useful selectivity. If the filter is a constant-Q design (constant percentage bandwidth), very high Q is required to separate a –90 dB second harmonic from its fundamental. This high Q will make it difficult to maintain constant, calibrated gain and thus good measurement accuracy across a frequency range. Furthermore, the constant-Q filter will have problems separating high-order harmonics from one another. For example, assume a 1 kHz fundamental frequency and a tunable constant-Q filter. When tuned to the tenth harmonic (10 kHz), the filter bandwidth

Figure 19. Nonlinearity symmetric around zero, resulting in odd harmonics.

is five times wider (in hertz) than when tuned to second harmonic, and it will be impossible to accurately measure a low-amplitude tenth harmonic surrounded by higher-amplitude ninth and eleventh harmonics.

Total Harmonic Distortion (THD)

THD (not to be confused with THD+N, Total Harmonic Distortion plus Noise, discussed below) is normally a computation from a series of individual harmonic amplitude measurements, rather than a single measurement in itself. THD is the square root of the sum of the squares of the individual harmonic amplitudes. THD specifications will commonly describe the highest order harmonic included in the computations; for example, "THD through 5th harmonic." THD is not a commonly used measurement, since it requires a fairly unusual type of analyzer to measure individual harmonics down to good performance levels, plus an automatic or manual computation on the results. Note that many older analyzers of the THD+N architecture are labeled THD on their panel and many people refer to a THD measurement when they really mean the THD+N technique.

Total Harmonic Distortion Plus Noise (THD+N)

By far the most common method of distortion measurement is the THD+N technique. Figure 21 is a simplified block diagram of a THD+N analyzer. The key functional block is a tunable notch filter (bandreject filter). In operation, this filter is manually or automatically tuned to the fundamental sine wave frequency, so that the fundamental is greatly attenuated. The filter is designed to have virtually no insertion loss at second harmonic and higher, so harmonics are passed essentially unattenuated. Wideband noise, ac mains-related hum, and any other interfering signals below and above the notch filter frequency are also passed unattenuated; hence the "+N" (plus noise) portion of the name.
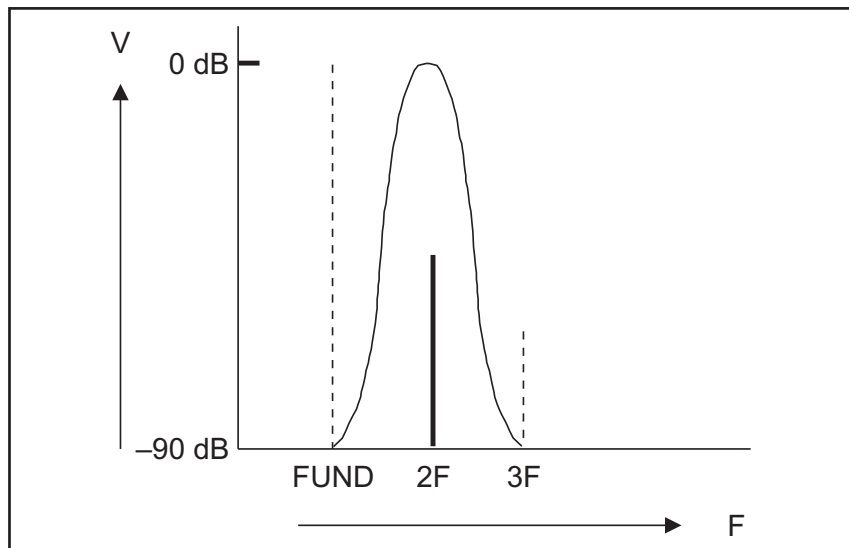
Figure 20. Bandpass filter for 90 dB rejection of fundamental.

There are several merits to the THD+N technique. From a circuit implementation technique, it is practical to design a notch filter which can attenuate the fundamental by 100 to 120 dB with only a few tenths of a dB insertion loss at second harmonic and above. From the standpoint of measurement concept, the THD+N technique is appealing since anything in the DUT output other than the pure test signal will degrade the measurement. A low measured THD+N value not only says that harmonic distortion is low, but also that hum, interfering signals, and wideband white noise are also at or below the measured value. Thus, THD+N comes closer than any other distortion measurement technique to providing a single figure of merit which says that there are no major problems in the DUT.

Figure 21. Block diagram, THD+N analyzer.

Noise Bandwidth for THD+N Measurements

The noise portion of a THD+N measurement may be negligible. If, for example, the wideband white noise and hum are 90 dB below the output level of a DUT and the second harmonic distortion is 1% (–40 dB), the noise has no practical effect on the reading and a THD+N measurement will exactly agree with a THD computation from individual harmonic measurements. If the DUT has quite low distortion, however, noise may be the limiting item. Most well-designed modern audio electronics devices such as equalizers, preamplifiers, and power amplifiers have distortion below their wide-band white noise when operated at levels below clipping. For such devices, it is critical that a measurement bandwidth be specified for THD+N. This implies that a THD+N analyzer then must have a selection of bandwidth-limiting filters so that measurements can be made with the specified noise bandwidth. Since most audio analyzers use many of the same circuit blocks for level, noise, and THD+N measurements, the same bandwidth-limiting filters are typically used for THD+N and for noise. Common values of high-frequency cutoff are 22 kHz, 30 kHz, and 80 kHz. Weighting filters are rarely used for THD+N measurements.

Interference Signal Effects on THD+N

Since THD+N is intrinsically a wide-band measurement, sensitive to everything between the band limits after removal of the test signal, it is susceptible to interfering signals. If an interfering signal exists in the DUT, it acts to establish a floor below which measurements cannot be made. Probably the most common type of interfering signal is ac mains (power line) hum. If an amplifier is being tested at a 1.00 V output level and 10 mV of hum voltage at 50 or 60 Hz exists, this will be measured as 1.0% distortion by a THD+N analyzer. Even if the true distortion of the amplifier is less than 0.01%, it will be masked by the constant 10 mV (1%) of hum voltage. The exception to this statement occurs when the test oscillator and the THD+N analyzer notch filter are tuned to the interfering signal frequency. The notch filter then attenuates the interfering signal in addition to the generator signal, and the measurement may approach the true distortion level of the amplifier. Figure 22 is a graph of THD+N versus frequency for a hum-limited electronic device. The curve is constant and horizontal all across the spectrum except at the 60 Hz mains fre-
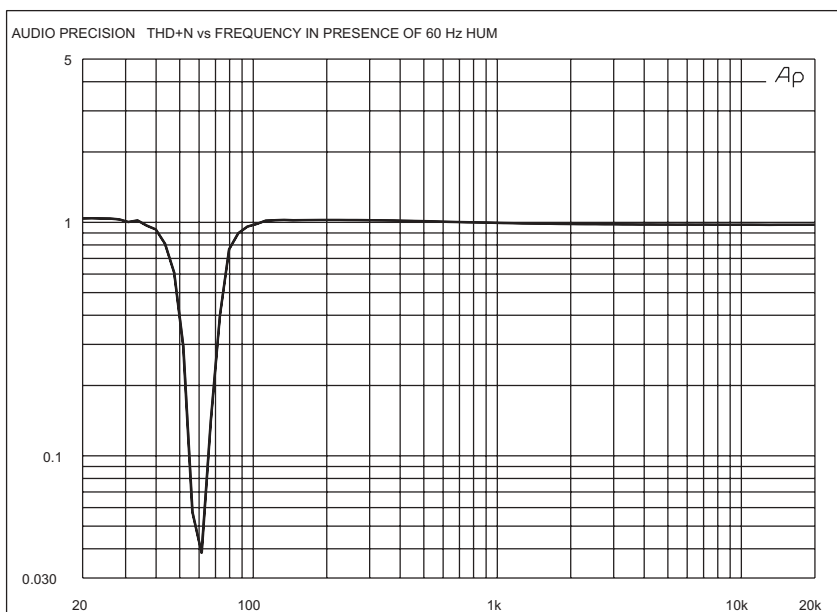
Figure 22. THD+N vs frequency in hum-limited device.

quency, where a sharp dip in THD+N occurs when the 60 Hz hum is attenuated by the analyzer notch filter. This characteristic THD+N vs frequency curve shape should always be recognized as an indication that a hum problem (or other interfering signal) exists which must be cured before accurate measurements can be made. Typically, this curve may result from a broken shield on one of the cables connecting the DUT to the test set. If spectrum analysis capability is available, it is good testing practice to make a spectrum analysis after connecting any new test setup to be sure that hum or other interfering signals are low enough to not cause limitations on THD+N and noise measurements.

Noise-Limited Devices

One may intuitively expect that distortion increases with increasing signal level in a device; turn up the volume of a cheap radio and it gets both louder and more distorted. Measurements will show that many well-designed audio electronics devices actually exhibit a decrease in THD+N percentage with increasing signal across most of the dynamic range of the device. Figure 23 is a typical graph of THD+N versus input signal amplitude on such a device. This curve does not conflict with the intuitive expectation that distortion should increase with amplitude. The device tested is in fact noise-limited across most of its dynamic range; the wideband noise contains more power than the harmonic distortion products. The output noise of the device is constant and independent of signal level. But, the vertical axis of the graph is THD+N expressed as a percentage of the signal amplitude and the signal amplitude increases linearly to

AUDIO PRECISION THD+N (% OF OUTPUT LEVEL) vs INPUT LEVEL
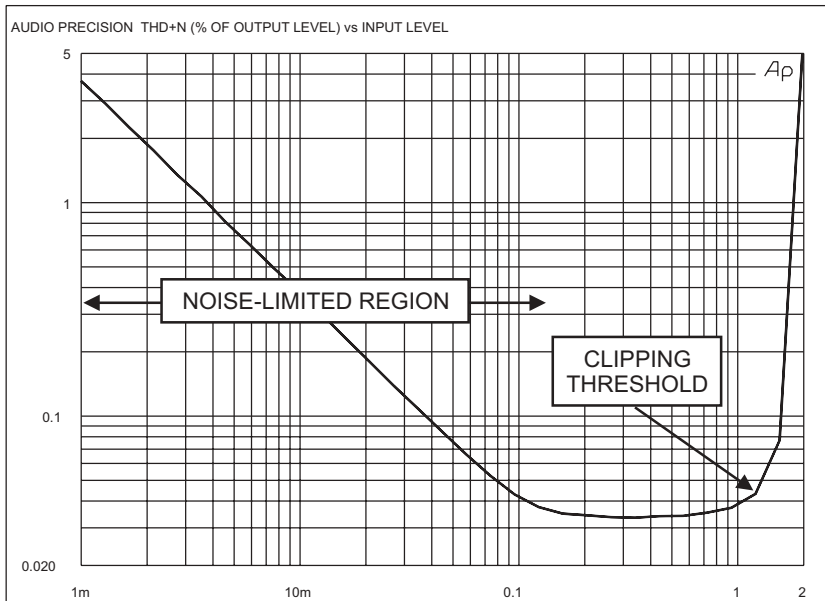
Figure 23. THD+N (in per cent) vs amplitude, noise-limited device.

the right. So, the constant output noise represents a decreasing percentage of the total output as the signal amplitude increases, and the THD+N curve thus falls to the right throughout the noise-limited portion of its operation. Narrow-band individual harmonic measurements (or THD computed from individual harmonic measurements) will show this characteristic to a much lesser degree, since the narrow noise bandwidth of the measuring instrument rejects most of the wideband noise.

Relative vs Absolute Distortion Units

Distortion is typically expressed relative to the signal level, either as percentage of the signal or as dB below the signal. Some THD+N analyzers also have the ability to display and graph THD+N in absolute units (V, dBV, dBu, etc.) instead of relative to the signal amplitude. Figure 24 is a graph of the same device over the same amplitude range as Figure 23, but expressing the THD+N in dBV (decibels below one volt). This presentation of the data makes it much clearer that the noise is constant and independent of signal level, and that ultimately actual distortion products rise above the noise level when the DUT is driven to higher levels and finally into clipping.

AUDIO PRECISION   THD+N (ABSOLUTE UNITS) vs INPUT LEVEL
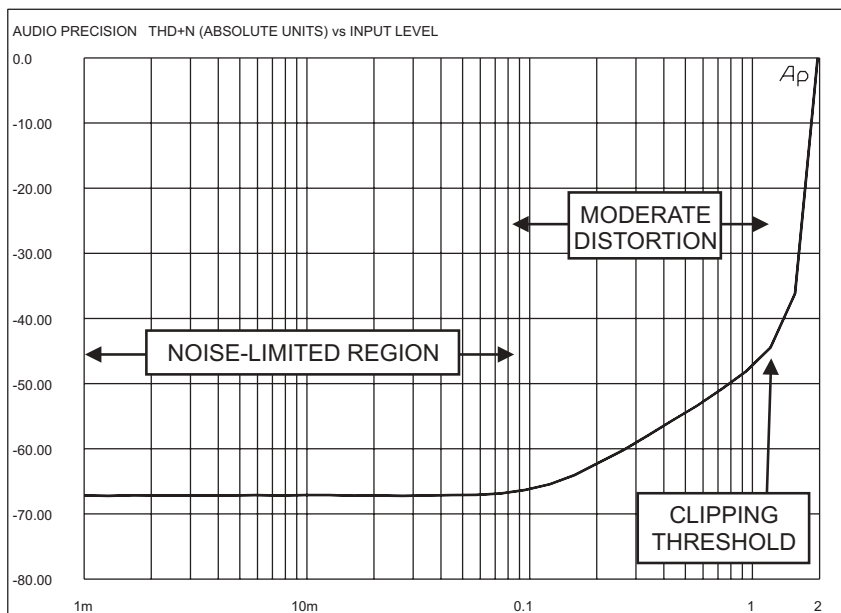
Figure 24. THD+N (absolute units), noise-limited device.

Harmonic Order vs Bandwidth; Limitations on Usefulness of Harmonic Distortion Techniques with Band-Limited Devices

The upper bandwidth limit serves as a limitation on the highest harmonic order which can contribute to a measurement. This harmonic order depends upon the fundamental frequency in use. For example, a 22 kHz band-limiting filter or a DUT intrinsically limited to 22 kHz will cause harmonics higher than the 22nd harmonic of a 1 kHz fundamental to be attenuated. This is not likely to affect distortion readings in most practical cases, since normally the lower order harmonics contain the majority of the harmonic energy. Only the second harmonic of a 10 kHz signal will pass the 22 kHz low-pass filter without attenuation, and all harmonics will be attenuated at fundamental frequencies above 11 kHz. For this reason, no harmonic distortion measurement technique (THD+N, computed THD, or individual harmonics) is a useful method of determining linearity of a band-limited device at frequencies above one-half the band limiting frequency. Since many non-linearities produce dominantly odd harmonics, it is more conservative to say that no harmonic distortion measurement technique is a reliable measurement of non-linearity at fundamental frequencies above one-third the band-limiting frequency. There are IMD (Intermodulation Distortion) measurement techniques which are quite useful in band-limited systems; see the IMD sections below. Many typical devices are intrinsically band-limited, including all digital audio devices, analog tape recorders, and broadcasting channels.

Instrumentation Requirements for THD+N Measurements

Important criteria for high-quality harmonic distortion measurement instruments include the following:

- The generator must have extremely low residual distortion. There is no practical method to separate distortion in the generator from distortion introduced in the device under test, so the generator distortion must be well below the lowest distortion value that is to be measured. Generally, a margin of 10 dB or greater is desired in order that the generator distortion not contribute to the reading.

- The residual distortion in the input stages and notch filter circuitry of the analyzer must also be lower than the lowest values to be measured.

- The analyzer input circuitry should have a high CMRR to reject common-mode noise.

- The analyzer must have a selection of band-limiting filters to control the noise bandwidth of the measurement.

- For ease and speed of operation, the instrument should have autoranging input stages, auto set level, automatically-tuning and nulling notch filters, and autoranging of the stages following the notch filter.

- The instrument detector must be true rms-responding for accuracy, since the signal measured often is largely noise or a complex combination of several harmonics.

Intermodulation Distortion (IMD)

All intermodulation distortion techniques use a stimulus signal more complex than a single sine wave. The most widely-used IMD techniques in professional audio, broadcast, and consumer audio use two sine waves as the stimulus signal. With two sine waves of frequencies F_1 and F_2 as the stimulus signal, a non-linear DUT output will consist of the original two sine waves plus an infinite number of IMD products given by the equation

$$m \cdot F_1 \mp n \cdot F_2$$

where m and n are all possible integers. The "order" of any particular IMD product is the sum of m and n. Thus, the order of some IMD products are as in the following table:

F_2-F_1	2nd order (even)
F_1+F_2	2nd order (even)
$2F_1-F_2$	3rd order (odd)
F_1-2F_2	3rd order (odd)
$2F_1+F_2$	3rd order (odd)
$3F_1-F_2$	4th order (even)
$3F_1+2F_2$	5th order (odd)

and so on.

The terms "odd" and "even" refer to whether the sum of ($m+n$) is an odd or even number. As in the case of harmonics, odd order IMD products are produced by symmetrical non-linearities in the device transfer function and even-order IMD products by non-symmetrical non-linearities.

A potential advantage of IMD measurements versus harmonic distortion measurements is that the test can be arranged so that many distortion products fall within the audio band, permitting linearity measurements of band-limited devices.

SMPTE/DIN IMD

The most common IMD measurement standards in the professional, broadcast, and consumer audio fields are the SMPTE and DIN methods. SMPTE (Society of Motion Picture and Television Engineers) standard RP120-1983 and DIN (Deutsches Institut fur Normung e.V.) standard 45403 are similar. Both specify a two-sine wave test signal consisting of a low-frequency high-amplitude tone linearly combined with a high-frequency tone at ¼ the amplitude (–12.04 dB) of the low frequency tone. The SMPTE specification calls for 60 Hz and 7 kHz as the two frequencies. The DIN specification permits several choices in both low and high frequency selection; 250 Hz and 8 kHz are a commonly-used set of frequencies which comply with the DIN specification. Other SMPTE-like or DIN-like signals are also sometimes used, such as 70 Hz and 7 kHz. Figure 25 shows such a signal in both time and frequency domain representations.

When such a two-tone test signal is fed to a non-linear device, intermodulation products appear as a family of sidebands around the high-frequency tone. The spacing between the high frequency tone and the first pair of sidebands (2nd-order sidebands, $F_2 \pm F_1$) is equal to the low frequency tone; the second pair of sidebands (3rd-order sidebands, $F_2 \pm 2F_1$) are spaced at twice the low frequency tone from the high frequency tone, etc. The percentage of
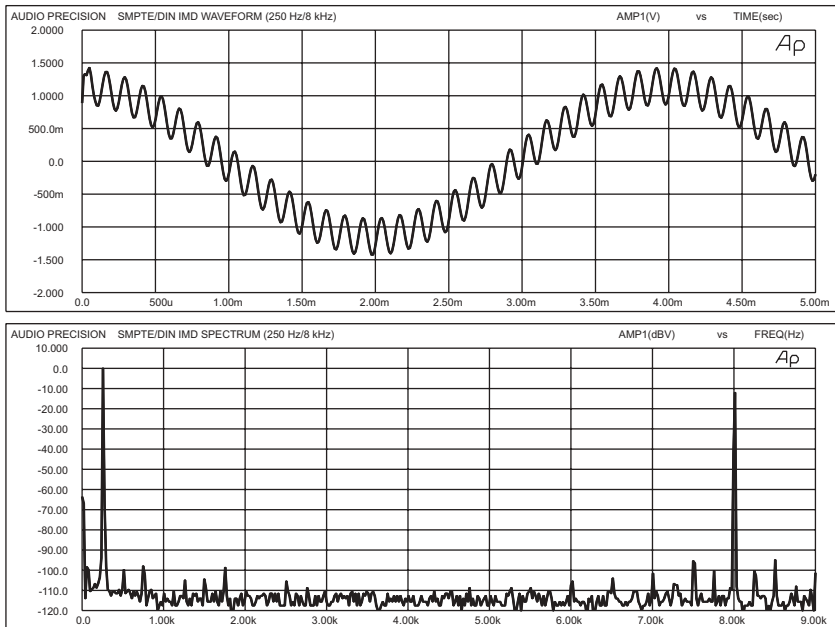
Figure 25. SMPTE/DIN IMD signal, time domain and frequency domain.

intermodulation distortion is defined as the percentage of amplitude modulation which these sidebands represent of the high frequency "carrier."

A typical SMPTE or DIN IMD analyzer has a simplified block diagram as shown in Figure 26. A high-pass filter first eliminates the low-frequency tone. The remaining signal is basically an amplitude-modulated signal, and is fed to an AM demodulator. The output of the demodulator consists of the sidebands translated down to baseband; for example, with a SMPTE test resulting in 2nd- and 3rd-order IMD products, the upper and lower sidebands would be translated down by the demodulator to 60 Hz and 120 Hz components. A low-pass filter follows the demodulator to remove any residual high-frequency "carrier," and the remaining signal is measured by a true rms detector. The detector output is calibrated in percentage amplitude modulation of the "carrier."

Figure 26. Block diagram, SMPTE/DIN analyzer.

SMPTE and DIN intermodulation testing has several advantages for testing audio devices. Like most intermodulation test signals, the more complex test signal is a step closer to simulating actual program material. The spectral balance between the higher-amplitude low frequency tone and lower-amplitude high frequency tone is also somewhat similar to music and voice spectral distribution. Many SMPTE/DIN IMD analyzers have a bandwidth after the demodulator of about 700 Hz, so at least second and third order IMD products will be measured if the low-frequency tone is below approximately 250 Hz. This makes the technique sensitive to both even and odd order distortion mechanisms. With a final noise bandwidth of about 700 Hz, the technique is substantially less sensitive to noise than THD+N with noise bandwidths of 20 kHz or greater. The use of a low frequency tone causes the IMD products to fall within a fairly narrow band around the high frequency tone, so the SMPTE and DIN techniques can be used to explore the linearity of band-limited systems quite close to their band-limiting frequency.

SMPTE/DIN Instrument Criteria

A high-quality SMPTE/DIN IMD test set should have very low residual intermodulation distortion in both the generator and analyzer; values of –90 dB (0.003%) or lower are readily available. The analyzer detector should be true rms responding, since the signal may contain several frequencies and thus is non-sinusoidal. It is desirable that both the low frequency tone and the high frequency tone can be varied or selected across a substantial range, to permit SMPTE-like or DIN-like tests on bandwidth-limited systems. For example, voice-bandwidth systems can be tested with a low frequency tone of 300 Hz to 500 Hz and a high frequency tone of 3 kHz if the generator and analyzer architecture permit. Similarly, a 20 kHz band-limited system can be tested by placing the high frequency tone at about 19 kHz, which provides a good test of high frequency linearity while still allowing several low-order IMD products to fall within the band limit. Some audio test sets make it possible to sweep the high frequency tone across the audio spectrum above some limit such as 2.5 kHz.

"CCIF," Twin-Tone, Difference-Tone IMD

The stimulus signal for this type of intermodulation distortion testing consists of two equal-amplitude high frequency signals spaced rather closely together in frequency. Common signals are 13 kHz and 14 kHz for 15 kHz band-limited systems, and 19 kHz and 20 kHz for systems with a full audio bandwidth. Figure 28 shows this type of IMD signal in both time and frequency domains. While such a stimulus signal will produce an infinite number of IMD products across a wide spectrum according to the basic IMD equations shown earlier, it is common in practical IMD analyzers to simplify the instrument architecture by measuring only the low-frequency second-order product falling at F_2 - F_1. If the maximum frequency spacing permitted between the two tones
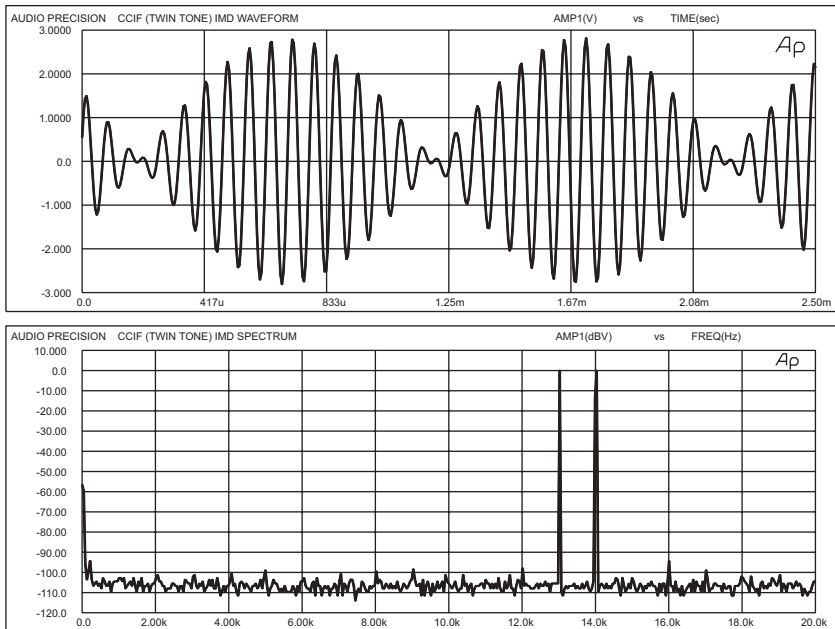
Figure 28. "CCIF" IMD signal, time domain and frequency domain.

is 1 kHz, the analyzer architecture can consist simply of a low-pass and/or bandpass filter at 1 kHz, followed by a voltmeter; see Figure 27.

This "twin tone" test permits stressing band-limited systems at their highest frequencies while still measuring an in-band IMD product. The most severe limitation of the simplified analysis technique is that it measures only the second order product. The simplified technique is thus not useful to measure distortion produced by non-linear transfer functions which are symmetrical about zero, such as the BH curves of magnetic tape; this method should not be used to measure tape recorders. The same twin-tone signal can be used with an FFT

Figure 27. Block diagram, CCIF IMD analyzer.

analyzer or other selective analyzer to measure IMD products of all orders and thus yield more information, but analysis of such information requires more skill than the simple "one number" readout of the simplified analysis technique.

"CCIF" IMD Instrument Criteria

Both the generator and analyzer should have low residual intermodulation distortion. With the simplified analysis technique which measures only the low frequency second order product at F_2-F_1, the analyzer will profit from extremely narrow bandwidth. Analyzers which use both a low-pass filter and a one-third octave band-pass filter at the difference frequency can provide noise bandwidths of 200 Hz and less, reducing the region of noise-limited measurements to quite low levels. It is possible by this technique to provide analyzers with residual difference-frequency IMD and noise below –100 dB. It is desirable to be able to set the generator tone frequencies to different spacings across a wide frequency range in order to comply with various versions of this testing technique. Some test sets make it possible to sweep the two-tone pair across a wide frequency range above approximately 4 kHz, maintaining a constant spacing between the two tones, so that the analyzer can continuously display the second-order product as a function of tone pair center frequency.

Figure 29. Dynamic/Transient IMD signal, time domain and frequency domain.

Dynamic/Transient Intermodulation Distortion (DIM/TIM)

Many audio workers have felt that some audible distortion mechanisms were triggered by changing program material, but not evident with steady-state stimulus such as a single sine wave. In particular, amplifiers with high amounts of negative feedback were singled out as a possible source of such problems due to the time delay inherent in negative feedback loops. The theory was that when a rapidly-changing signal was fed to such an amplifier, a finite time was required for the correction signal to travel back through the feedback loop to the input stage and that the amplifier could be distorting seriously during this time. Dynamic and transient intermodulation test techniques were developed in an effort to isolate these phenomena.

Most proposed DIM/TIM test techniques therefore implement a signal with a rapidly-changing (high slew rate) component. The most popular technique was proposed by Schrock and Otala. The signal (see Figure 29) consists of a band-limited square wave (typically around 3 kHz) plus a high frequency sine wave "probe tone" of one-fourth the peak-to-peak amplitude of the square wave. The rise and fall sections of the square wave stress any portions of the
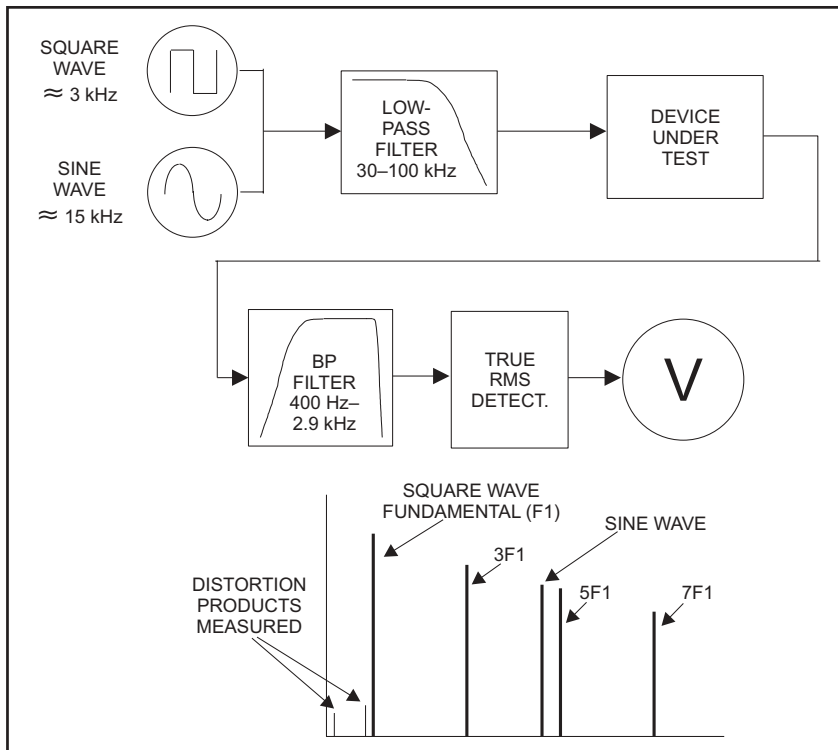
Figure 30. Block diagram, simplified DIM/TIM analyzer.

circuit which have slew rate limitations, and intermodulation products of the sine wave with the fundamental and various harmonics of the square wave will indicate problems. Fully examining the spectrum of such a signal requires an FFT analyzer or other selective spectrum analyzer and a degree of skill. Skritek proposed a simplification in which, by proper selection of sine wave frequency, both an even order and an odd order IMD product would fall into the mid-frequency spectral region below the fundamental of the square wave. This simplified technique can be rather simply implemented with a sophisticated bandpass filter followed by an rms detector, and provides a "one number" readout which does not require operator skill to interpret. Figure 30 is a block diagram of this instrument architecture.

DIM/TIM Instrument Criteria

The square wave generator portion of the oscillator must be followed by selectable low-pass filtering, since both 30 kHz and 100 kHz band limited versions of the signal are proposed by Otala. The stress which the test signal places on the device under test is principally determined by this band limit value. The analyzer must have an extremely well-designed bandpass filter (particularly the lowpass section) in order to reject the square wave fundamental component while providing a low residual distortion value for IMD products falling only a few hundred Hz below the square wave fundamental. Two versions of the technique have been proposed; one for relatively non-band-limited devices such as power amplifiers, and one for band-limited devices or systems such as 15 kHz broadcast links. In order to comply with both versions, the square wave frequency must be settable to either 3.15 kHz or 2.96 kHz and the sine wave probe tone to either 15 kHz or 14 kHz. The detector must be true rms since it is typically measuring a complex signal consisting of at least two components at different frequencies plus noise.

Other IMD Techniques

Several other techniques have been proposed for IMD measurement by workers including Cordell and Thiele. Some use three sine waves; some use a sine wave plus a ramp (triangle) waveform. Few have yet been designed into commercial instrumentation. There are also other IMD measurements which are used in audio-related fields such as telecommunications, but not in consumer audio, pro audio, or broadcasting.

Frequency Measurements

Most modern audio test sets include frequency measurement capability. This capability is almost invariably implemented by standard frequency counter architecture shown as diagrams A and B in Figure 31. In A, a gate is opened for a precise period of time determined by a quartz clock and cycles of
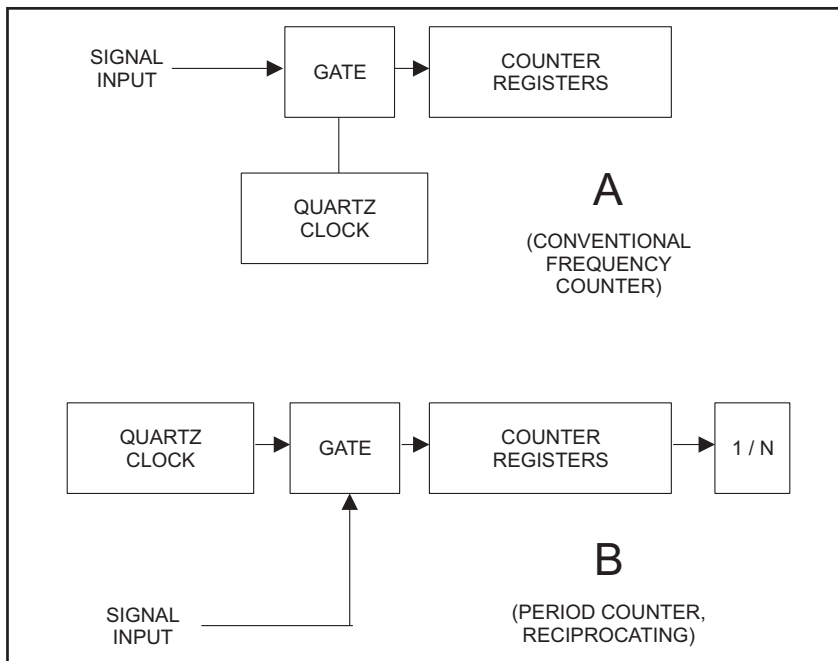
Figure 31. Block diagrams, conventional and period/reciprocal frequency counters.

the signal are accumulated into counter registers while the gate is open. If the gate is open precisely 1 s, the accumulated count is the signal frequency in cycles per second, or hertz (Hz). Other values of precisely-known gate time may be used if the counter register is appropriately scaled before display. Longer gate times are required in order to obtain better resolution.

In B, a "squared-up" version of the signal is used to control the gate and pulses from a quartz-based clock are accumulated into the counter. This count is proportional to the signal period; taking the reciprocal of the count provides a value proportional to the signal frequency, and scaling for the clock frequency provides the exact signal frequency. Whenever the signal frequency is lower than the clock frequency, this "period and reciprocate" technique of B always provides better resolution for any given measurement interval than the conventional counter architecture A. Since typical quartz clock frequencies are 10 MHz or higher, this is the preferred technique when measuring audio frequencies. It is also practical to hold the gate open for several or many cycles of the signal waveform in order to accumulate still more clock cycles into the register, if the register count is then divided by the number of cycles of waveform for which the clock was open. This period-averaging reciprocal counting technique is the most valuable for audio. It can permit selection of a desired measurement interval and then will provide maximum possible resolution and trigger noise rejection for the particular clock frequency and measurement interval in use.

Frequency Counter Instrument Criteria

Sensitivity of the input and triggering circuitry should permit frequency measurements on relatively low-amplitude signals, preferably down into the tens of millivolts or even lower. The reciprocal, period-averaging technique is by far the most useful at audio. Frequency accuracy will be proportional to the accuracy and stability of the quartz crystal oscillator, but actual audio applications can normally be satisfied with quartz oscillators far below the state of the art. Quartz oscillators with specifications within a few parts per million are readily available, but accuracies greater than a few hundred parts per million are rarely required for most audio work.

The basic frequency counter architecture has no frequency selectivity and will normally count the highest amplitude signal present or give false readings if the signal is complex. For certain audio applications, it is desirable to be able to pass the signal through a bandpass filter before connecting it to the frequency counter, allowing independent measurement of the frequencies of several components of moderately-complex signals such as IMD test signals.

Phase Measurement

Measurement of the phase difference between two audio signals of the same frequency is commonly required. Adjustment of the azimuth angle of stereo or multi-track analog tape recorder heads is best done by measuring the phase difference between the outside tracks and adjusting azimuth for zero phase. Stereo transmission paths must have matching phase characteristics if they are to properly preserve the intended stereo effect. Input-to-output phase shift of an amplifier is an important characteristic when designing feedback circuits. Phase is deliberately manipulated in several types of audio processing equipment.
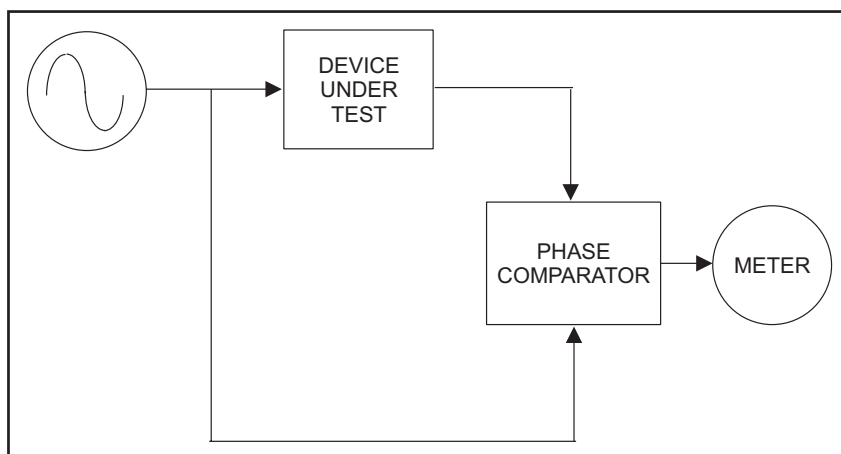
Figure 32. Input/output phase measurement.

Phase may be measured, or at least estimated, with an oscilloscope which permits one signal to be fed to the vertical and the second signal to the horizontal deflection inputs. The resulting Lissajous pattern will vary from a circle through ellipses to diagonal straight lines with positive or negative slope, depending on the particular phase relationship. The Lissajous technique provides its best accuracy and resolution at or near zero and 180 degrees, and relatively poor resolution at 90 and 270 degrees.

More precise techniques which can provide digital numeric output typically operate by comparing the zero-crossing time of the two signals. These techniques can yield constant resolution and accuracy at all values of phase difference.

Input-Output Phase

Phase meters have two sets of input terminals and display the phase difference between signals of the same frequency at the two sets of terminals. When the phase meter is integrated into a general-purpose two-channel (stereo) analyzer, the phase meter normally measures the phase difference between the two channels.

It is often desirable to measure the phase difference between the input and output of a device. To do this with an integrated phase meter, connect the DUT output to one set of analyzer input terminals. Connect the other analyzer input across the connection from the generator output to the DUT input terminals (see Figure 32). On sophisticated analyzers, this capability may be available via internal cables and switches, via a generator monitor path, which also manages any necessary polarity inversions in order to obtain the correct display.

Phase Meter Instrument Criteria

Phase resolution, accuracy, the minimum amplitude at which phase may be measured, and the acceptable frequency range of signals are the most important criteria for phase meters. Resolution of 0.1 degrees and accuracy of 1 to 2 degrees are readily available, and anything less will not be useful in determining phase match of two transmission paths in order to determine whether they are suitable for stereo transmission. As with frequency measurements, it is quite desirable to be able to measure signals with amplitudes down into the tens of millivolts.

If the zero-crossing technique is used, it should measure both positive-going and negative-going zero-crossings of the signal and use the average of the two in order to provide accurate phase measurements of signals whose duty cycle is not precisely 50%.

Wow and Flutter Measurements

Wow and flutter is the name used to describe signal degradation caused by short-term speed variations of analog recording and playback devices such as tape recorders and turntables. If the speed varies during either recording or playback, the result is a momentary change in pitch (frequency) of the signal. Thus, wow and flutter is actually frequency modulation (FM). Such variations are commonly caused by mechanical imperfections in rotating parts such as capstans, idlers, and pulleys, or short-term variations in the speed of the motors operating the mechanism. Wow is normally used to describe speed variations at a low rate, perhaps below 6 Hz. Flutter describes speed variations above about 6 Hz. Low rate variations are caused by defects in rotating parts of large circumference, and flutter by rotating parts of smaller circumference. An idler with a flat spot rotating ten times per second at a given tape speed will produce a flutter component at 10 Hz. The frequency of wow and flutter components normally scales up and down with changes in the tape speed of tape machines.
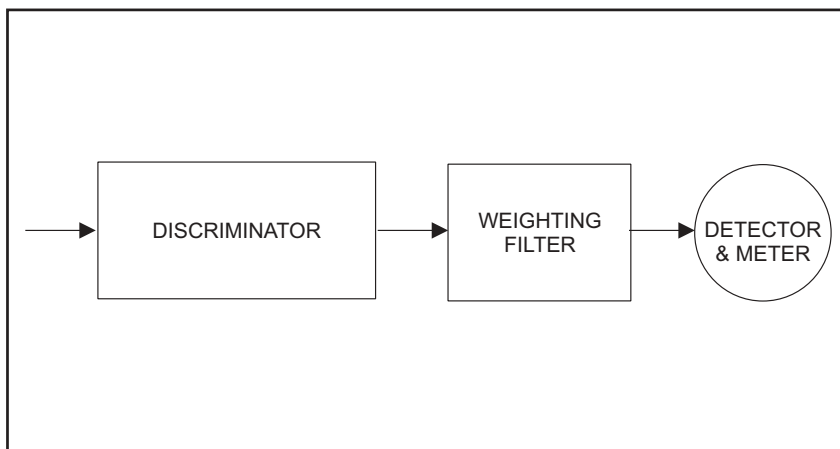
Figure 33. Block diagram, wow and flutter analyzer.

When a pure sine wave signal is recorded and reproduced from a system whose pitch depends upon speed accuracy, a measurement of the frequency modulation of the reproduced signal is a measurement of wow and flutter. Peak wow and flutter in percent is defined as the peak frequency deviation stated as a percentage of the recorded signal frequency ("carrier frequency"). Test tones of either 3.00 kHz or 3.15 kHz are normally used for the IEC/DIN, NAB, and JIS standard techniques of wow and flutter measurement. Thus, a wow and flutter value of 0.1% when using a 3 kHz test tone would correspond to 3 Hz peak deviation of the FM signal. A wow and flutter meter in principle consists of an FM discriminator which converts frequency deviations of the 3 kHz or 3.15 kHz "carrier" into proportional amplitude variations which may be measured with an ac voltmeter; see Figure 33. Since wow is defined below

0.5 Hz, the voltmeter must have extended low frequency response compared to that necessary for most audio measurements.

The IEC/DIN, NAB, and JIS standards vary in the detector response and time constants used following the FM discriminator. The table below shows the recommended test tone frequency and detector response of each of these standards. Test tone frequency is not normally critical, however; in practice, the same percentage numbers will be obtained at either 3 kHz or 3.15 kHz.

Standard	Frequency	Detector	Time
IEC/DIN	3.15 kHz	peak	normal
NAB	3.0 kHz	rms	normal
JIS	3.0 kHz	rms	slow

The IEC/DIN standard uses peak-reading detectors and thus provides the highest readings for any given conditions. The NAB and JIS standards are rms-calibrated and yield lower numbers. The JIS standard has very long time constants after the detector, thus averaging out short-term variations in flutter.

A weighting filter response has been defined which can be inserted between the discriminator output and the voltmeter. The shape of the weighting filter follows the perceived irritation factor to human listeners of wow and flutter at

Figure 34. Frequency response, wow and flutter weighting filter.

different rates. Figure 34 shows the weighting factor response, which peaks at about 4 Hz. The same filter is specified in all standards.

Measurement Technique

Wow and flutter is normally measured by use of a standard reference tape or disk which was recorded on a machine with very low wow and flutter. This tape or disk is played while the measurement is made. On tape machines, measurements should be made at various points on the reel or cassette including the beginning, middle, and end of the tape since the wow and flutter is likely to vary considerably under these conditions. Tape speed error can be measured at the same time as wow and flutter, since the reference tape was recorded on a machine with precisely-adjusted speed. Therefore, any deviation in the average reproduced frequency from the frequency specified in the reference tape documentation is a measurement of the tape machine deviation from nominal speed.

On very high-quality professional tape machines, it is possible that purchased standard reference tapes may contain more wow and flutter than the machine being tested. This can be checked by recording a 3 kHz or 3.15 kHz test tone on blank tape and then rewinding and reproducing it. If the wow and flutter measurements are lower under these conditions than when using a commercial reference tape, the self-recorded tape should be used. However, wow and flutter should never be measured during the simultaneous record-reproduce mode which is possible with a three-head tape recorder. The value of time delay between record head and reproduce head at any given tape speed causes a comb filter effect in the measurement. The measurement will be completely insensitive to flutter at flutter frequencies falling exactly at the comb filter frequencies, which may cause erroneously-low readings. Instead, the tape should be recorded, rewound, and then played for measurements.

Instantaneous vs Processed Readings

When a wow and flutter meter is observed during operation of a turntable or tape machine, the indication will be quite variable. An operator will develop an impression of the average or maximum value of flutter, but that impression could vary between operators. Graphs can be made on chart recorders or computer screens of the variations of wow and flutter versus time. Often, it is desired to have a "single number" for wow and flutter which can be entered into a specification sheet. In this case, it is common to use the "two sigma" value. The term "sigma" implies statistical standard deviation, but that is not how commercial wow and flutter meters actually work in their "two sigma" modes. Statistical analysis of a variable requires obtaining a set (of sufficient size to be representative) of sample measurements and computing the standard deviation of the set. If the probability distribution of the variable matches the Gaussian ("bell-shaped") probability distribution, then approximately 5% of the total

AUDIO PRECISION   W&F PEAK & 2-SIGMA vs TIME, CASSETTE DECK

Figure 35. Instantaneous vs 2-Sigma wow and flutter vs time.

samples will be more than two standard deviations (two sigma) above the mean value and 95% will be below this "plus two sigma" value. However, there is no reason to assume that variations of wow and flutter with time will follow a Gaussian probability distribution. Furthermore, true statistical analysis requires collecting the entire set of data before making the analysis, which is inconsistent with a quasi-real-time instrument. Therefore, the technique historically developed and used in all commercial wow and flutter meters is to define the value exceeded only 5% of the time as the "two-sigma" value. Both analog and digital techniques have been developed which provide this two sigma value as a running indication. Figure 35 shows the results of unprocessed, short-term measurements of wow and flutter versus time over a 30-second period plus the computed two-sigma result of these measurements.

Wideband (Scrape) Flutter

The IEC/DIN, NAB, and JIS techniques all use 3 kHz or 3.15 kHz test tones and a post-discriminator bandwidth of 200 Hz. They are designed for sensitivity to rotationally-induced imperfections of the machines. There are also other modulation effects in tape machines involved with bending of the tape as it wraps around capstans, idlers, and the tape heads, and with tape stretching effects. These mechanisms generate frequency modulation of the recorded signal across a much wider spectrum, up to approximately 5 kHz. Testing for such problems without aliasing effects requires use of a higher- frequency "carrier" signal. Manquen developed "scrape flutter" and "high band flutter" measurement techniques using a 12.5 kHz test tone. Measurement of

FM effects with a 5 kHz bandwidth at a 12.5 kHz test tone requires recorder bandwidth to at least 17.5 kHz, so this method is limited to wide bandwidth machines. Some professional tape machines have incorporated adjustable idlers specifically to reduce the scrape flutter effect, and proper adjustment of such idlers requires access to a scrape flutter meter.

Spectrum Analysis of Wow and Flutter

Wow and flutter due to imperfect rotating parts will have a frequency (modulation rate) determined by the rotational speed of the part. An idler with a bump rotating five times per second at a particular tape speed will produce a 5 Hz frequency modulation of the test tone. If an FFT or other high-resolution spectrum analysis is made of the discriminator output while measuring wow and flutter, the peaks in the spectrum analysis may be used to determine which rotating parts are causing the problem. The circumference of the part may be calculated from the frequency of the spectral peak and the known tape speed. Figure 36 shows a spectral peak at 7.5 Hz. At a tape speed of 7.5 ips, this correlates to a pulley with a circumference of exactly one inch.

*Figure 36. FFT spectrum analysis of wow and flutter.*

Measurements With Noise Stimulus

A sine wave is not the only useful type of signal for certain types of audio measurements. Frequency response can be measured with a noise stimulus. A basic potential advantage of noise stimulus is that energy is present at all frequencies simultaneously, so measurements can be more rapid than waiting for

a sine wave to sweep or step through the spectrum. Either white noise or pink noise can be an appropriate signal, depending upon the characteristics of the analyzer to be used. Both random and pseudorandom noise signals can be used, each with its own advantages and limitations.

White Noise

White noise is the term applied to noise which is distributed across the frequency spectrum in such a manner that there is equal power per unit bandwidth. For example, with white noise the 20 Hz bandwidth between 20 Hz and 40 Hz would contain the same noise power as the 20 Hz span between 10,000 Hz and 10,020 Hz. Thermal noise produced in active components tends to be approximately white in distribution. A graph of noise power versus frequency will be flat as a constant-bandwidth bandpass filter is scanned across the spectrum of a white noise signal.

Pink Noise

Pink noise is used to describe noise which has constant power per percent bandwidth, per octave, per decade, etc. With pink noise, the octave between 20 Hz and 40 Hz would contain the same noise power as the octave between 10,000 Hz and 20,000 Hz. A graph of noise power versus frequency will be flat as a constant-percentage-bandwidth (constant Q) bandpass filter is scanned across a pink noise spectrum. Pink noise has a "mellow, balanced" sound to the human ear compared to white noise, which sounds brilliant and "hissy" due to the additional energy at high frequencies. The energy versus frequency distribution of voice and music is more similar to pink noise than it is to white noise.

Analysis of White Noise

White noise, with constant power per unit bandwidth, is normally analyzed by a selective analyzer with constant bandwidth at any frequency. Scanning, superheterodyne spectrum analyzers and FFT (Fast Fourier Transform) analyzers are examples of constant bandwidth analyzers. The combination of a white noise signal and a constant bandwidth analyzer will produce the same frequency response graphs of a device that a swept sine wave and wideband level meter will produce, except for frequency resolution detail (see below).

Analysis of Pink Noise

Pink noise is normally analyzed by a selective analyzer with constant percentage bandwidth at all frequencies. Examples of such analyzers are the Real Time Analyzers (RTAs) often used in acoustical work, or an analyzer which can sweep or step a constant-Q filter across the spectrum. An RTA essentially consists of a bank of filters of equal Q, each tuned to progressively higher fre-
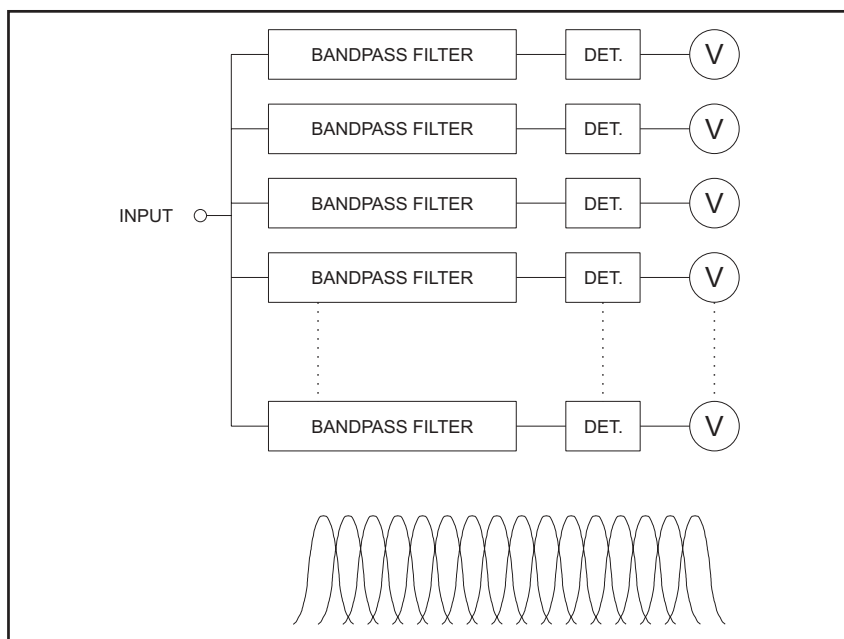
Figure 37. Block diagram, real time analyzer (RTA).

quencies across the band (see Figure 37). Each filter is followed by a detector and indicator. The input signal is split out to all filters. The number of filters and their center frequencies are chosen to be consistent with the bandwidth of each filter, so that adjacent filter curves cross at the –3 dB point. To cover the 20 Hz–20 kHz spectrum with 1/3 octave bandwidth filters in an RTA, for example, requires about 31 filters. Using ½ octave bandwidth filters permits the same spectrum to be covered with 20 filters. The combination of pink noise and a constant-Q analyzer will produce the same response graphs that the swept sinewave technique produces, except for detail.

Frequency Resolution Detail in Noise Testing

Since noise energy extends across the entire spectrum and the selective analysis filters in the analyzer are not infinitesimally narrow, one can expect poorer resolution from noise-stimulus testing than sine wave-stimulus testing. If the device under test is relatively flat with no areas of rapid response change and no sharp dips or peaks, noise and sine wave methods will yield identical results. If the DUT response has a sharp notch when measured with a swept sine wave, the notch will appear much more shallow and rounded with noise stimulus since the bandwidth of the analyzer filter modifies the measured shape. This reduction in frequency resolution would be considered a disadvantage when measuring devices such as parametric equalizers. However, some audio workers consider this lower resolution an advantage when measuring acoustical devices such as loudspeakers, particularly if the goal is to obtain results

which correlate well with human perceptions. Psychoacoustic experiments have shown that the effective selectivity of the human ear and brain is similar to a nearly-constant-Q filter with a bandwidth of about 1/5 octave at low frequencies, broadening slightly to about 1/3 octave at mid and high frequencies. The argument is thus advanced that resolution greater than 1/3 to 1/5 octave in loudspeaker measurements shows detail that humans cannot hear. However, this argument is not universally accepted.

Pseudorandom vs Random Noise Signals

A random noise signal is one whose amplitude versus time profile never repeats. A pseudorandom noise signal appears to be random when observed over a limited range of time, but is seen to exactly repeat over and over when examined over a long time period. It can be shown that a true random noise signal, observed over a sufficiently long time, contains energy at all frequencies. A pseudorandom noise signal contains energy only at discrete frequency points in the spectrum. These discrete points correspond exactly to the repetition rate of the pseudorandom sequence and all harmonics of that frequency. For example, if a pseudorandom signal sequence repeats exactly 3.8 times per second, the spectrum of this signal consists of energy at exactly 3.8 Hz, 7.6 Hz, 11.4 Hz, etc. See Figure 38 for a high-resolution spectrum analysis of energy from such a pseudorandom noise signal.

The advantage of pseudorandom noise comes from synchronizing the analysis device with the pseudorandom repetition cycle. If an analyzer meter with an integration time of exactly 250 milliseconds, for example, is used to measure a pseudorandom noise signal with a 250 millisecond repetition period, the analyzer reading will be perfectly stable. The erratic, varying, "peaky" meter readings normally associated with random noise measurements are eliminated when pseudorandom noise and a matching-measurement-cycle length analyzer are used together.

The disadvantage of pseudorandom noise results from the fact that all its energy falls at specific, harmonically-related frequencies rather than as a continuous spectrum. If the analyzer selective bandwidth is sufficiently wide so that it includes a large number of adjacent harmonics of the basic noise repetition rate within its passband, the resulting detected output energy will be essentially constant as the filter is tuned across the spectrum. If the analyzer bandwidth is small compared to the frequency spacing of the pseudorandom signal, however, the output of the analyzer detector will cycle up and down as the analyzer is tuned from a spectral peak through the valley between peaks and then onto the next peak. Figure 39 shows a 1/3 octave bandpass filter sweep of the spectrum from 10 Hz to 50 Hz with the pseudorandom noise signal shown in the previous figure. The 1/3 octave filter (constant percentage bandwidth of 20% of center frequency) has a –3 dB bandwidth of 2 Hz at a center frequency of 10 Hz, so it separates the spectral lines spaced 3.8 Hz apart. By the time the filter has tuned to a center frequency of 50 Hz, its –3 dB bandwidth is 10 Hz

Figure 38. High-resolution FFT, pseudorandom noise.

Figure 39. 1/3 octave spectrum analysis of pseudorandom noise.

wide and it thus encompasses several adjacent spectral lines, so little ripple is visible as the filter is swept.

Sets of Measurements

Many common audio tests are actually sets of measurements rather than single measurements. Examples include frequency response measurements, signal-to-noise ratio measurements, curves of THD+N versus frequency or amplitude, input-output linearity curves, etc.

Figure 40. Typical frequency response graph.

Frequency response, which is probably the most important and the most common audio measurement, normally consists of a series of level measurements, each at a different frequency. Even though each of the level measurements may be an absolute measurement (in V, dBu, dBV, etc.), the important information in a frequency response measurement is the relationship between the individual measurements. Frequency response is commonly plotted graphically, with some mid-band measurement (typically at 1 kHz) used as the reference. All other points are then plotted as variations from the mid-band reference level, and response will commonly be stated as "Fx dB over the range from y Hz to z kHz"; see Figure 40. Amplitude variations are normally plotted on the vertical axis versus test frequencies on the horizontal axis.

Frequency response measurements are greatly simplified by instruments which have these specific features:

AUDIO PRECISION   THD+N vs FREQUENCY, TYPICAL

Figure 41. Typical graph, THD+N vs frequency.

- A relative dB unit, with the ability to set the 0 dB reference to the level of the incoming signal. This setting is normally done while the mid-band reference stimulus frequency is present.

- Sweep capability for the generator

- Graphical results display capability, so that the response is plotted automatically instead of requiring the operator to manually plot a large number of measurements on graph paper.

It is also common to measure distortion using a number of different stimulus frequencies across the audio band and to plot the results on a scale with measured distortion plotted vertically against frequency horizontally; see Figure 41.

Measurement Dynamics and Reading Settling

Most audio measurements are taken under steady-state conditions. When stimulus is first connected, or when stimulus amplitude or frequency is changed, transients are produced in the audio test generator, the device under test, and the audio analyzer. If an operator is using manually-controlled and manually-read instruments, it is intuitive to wait until the indicated measurement has stabilized before noting the exact reading. With automatic audio test systems, features must be deliberately built into instrument hardware or software to assure stabilized results.

The most common solution to the problem of obtaining settled measurements is to simply introduce a time delay between any change in generator stimulus and acquisition of the resulting reading. The instrument designer or test program writer makes an estimate of the worst-case settling time and inserts this value as a fixed delay. While conceptually simple, this is not an efficient practice. Circuits normally settle faster at high frequencies than at low frequencies. Measurements made with a wide-band meter settle more quickly than selective measurements. If the analyzer includes autoranging, measurements settle much more quickly if the new value lies in the same level range as the previous value, avoiding the need to change ranges. Under noisy conditions, the measured value never settles to the same degree that a signal measurement will under good signal-to-noise ratio conditions. If the signal being measured is near in amplitude and very near in frequency to an interfering signal such as ac mains-related hum, a slow beat effect may result which will never settle. Attempting to allow for many of these conditions with a fixed time delay will produce very slow measurement times under other conditions, and the fixed delay will not cope with several of these problems at all.

Some of the more sophisticated software-based measurement systems include algorithms which mimic the processes which an experienced test engineer would go through in order to get the best possible data under any conditions. These algorithms are normally adaptive to the measurement conditions. Ideally, the measurement hardware will have a selection of response times, automatically selecting the fastest time which will produce specified accuracy at the present signal frequency. The software may then compare several or many successive samples of the stream of data, compute the difference between them, and finally accept as sufficiently settled a reading when the sample-to-sample difference has dropped below a threshold value selected by the test engineer when he sets up a particular test. This threshold value should be a relatively large value if a noise measurement is being made, for example, and a relatively small value when measurements at the limiting accuracy and resolution of the system are desired. "Safety valves" are also needed to cover the never-settled cases such as the beat with ac mains hum described above. Typically, this will be a time-out routine which will supply an average of the last several readings if no settled measurement occurs within some maximum "timeout" value. A sophisticated settling algorithm will also allow for averaging of a stream of readings when an integrated noise measurement is desired, and for disabling the settling algorithm when instantaneous values are required for measurements such as wow and flutter or phase jitter.

INPUT R (OHMS) vs DECREASE IN 600 OHM GENERATOR OUTPUT (dB) FROM UNLOADED
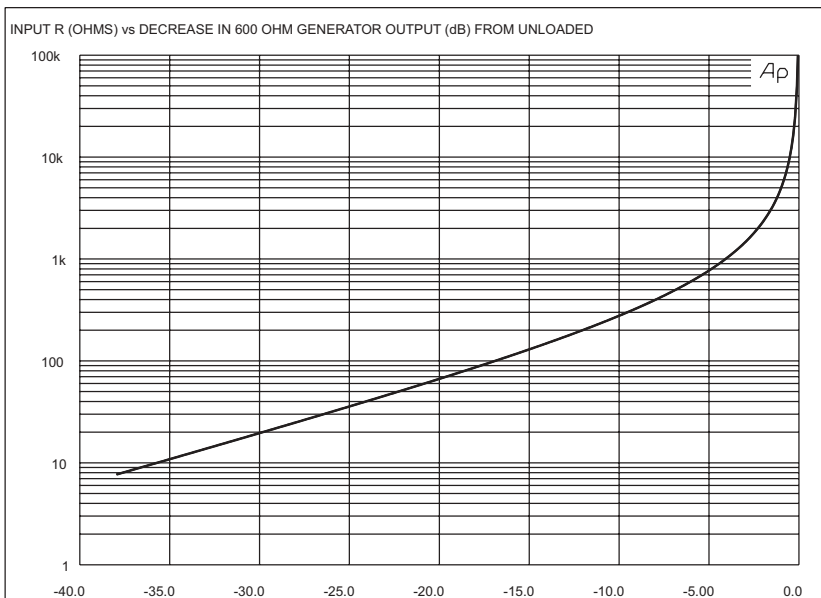
Figure 42. Input impedance vs voltage drop caused by connecting to a 600 Ω generator.

INPUT R (OHMS) vs DECREASE IN 600 OHM GENERATOR OUTPUT (dB) FROM UNLOADED

*Figure 43. Input impedance vs voltage drop caused by connecting to a 600 Ω generator, expanded at small values of voltage drop.*

Estimating Input Impedance

It is often necessary to measure, or at least make a reasonably accurate estimate of, the input impedance of a device. Consumer audio electronic equipment normally has quite high input impedances, on the order of 100 kΩ except for phonograph preamplifier inputs which are standardized at 47 kΩ (lower for moving coil pickups). Professional audio equipment input impedances are likely to be on the order of 10 kΩ and up for "bridging" inputs at line level, perhaps 1.5 kΩ for microphone inputs, and 600 Ω for terminated inputs. A good estimate of input impedances up through several thousand ohms can be made by noting the drop in output voltage of a 600 Ω generator when it is connected to the input in question. Figures 42 and 43 are graphs of voltage drop (in decibels) when the 600 Ω generator is connected to the circuit. An input impedance of 600 Ω will cause the generator output to drop exactly in half (–6.02 dB). This technique loses resolution as the device input impedance becomes high, but in practice it is often only necessary to know that the input is high and therefore will not significantly load the typical low output impedances of both professional and consumer electronic devices.

Output Impedance Estimation

Measuring, or estimating, the output impedance of audio devices is normally accomplished by removing all load impedances from the device (other than the audio analyzer input, typically 100,000 Ω), driving the device with a sine wave in order to produce an unloaded reference output level, and then connecting a known load resistance across the output and noting the voltage drop. The selectable 600 Ω input impedance available on some audio analyzers is a convenient value for estimating output impedances across the range from a few tens of ohms up through a few thousand ohms. Figures 44 and 45 are graphs of reduction of output level, in decibels, versus device output impedance when a 600 Ω load is connected across the device. Resistors of other known values may be used instead of 600 Ω to obtain the most accurate measurement; the output voltage of any device will drop exactly in half (–6.02 dB) when a load equal to the output impedance is connected across the device output.

OUTPUT R (OHMS) vs DECREASE IN DEVICE OUTPUT (dB) WITH 600 OHM LOAD

Figure 44. Output impedance vs voltage drop caused by loading with 600 Ω.

OUTPUT R (OHMS) vs DECREASE IN DEVICE OUTPUT (dB) WITH 600 OHM LOAD


Figure 45. Output impedance vs voltage drop caused by loading with 600 Ω, expanded at small values of voltage drop.

Characterizing a "Black Box"

Full documentation on the expected and specified performance of an audio device is not always available. Yet, it may be important to determine just what the device is doing in order to infer whether it is doing it correctly. Therefore, it should be instructive to consider what tests, and in what sequence, might be performed on an unknown "black box" in order to define its function.

Several generalizations can be made. First, very few devices will be damaged by feeding a signal which is too low in amplitude, but many can be damaged by too large a signal. Second, in both the professional-broadcasting equipment area and in consumer audio equipment, almost all devices are likely to include the output amplitude range around 1 V or 0 dBu (0.7746 V) in their range of specified and safe operation. The most fragile audio devices are normally loudspeakers, especially high-frequency devices (tweeters). Amplitude-modulated transmitters are also relatively fragile, though typically transmitters are equipped with protective circuits which will disable them at high modulation levels before permanent damage takes place, while most loudspeakers are not. So, a good first step with an unknown audio device is to connect an audio level meter across the output, connect a generator set to a midband frequency such as 400 Hz or 1 kHz at a low output level such as 1 mV or –60 dBu to the input, and quickly determine whether the output level exceeds about 1 V or 0 dBu. If the output level is below this range, the generator level can then be increased in steps of 6 to 10 dB while watching the output level until the output


Figure 46. THD+N and output level vs generator amplitude.

AUDIO PRECISION THD+N (%) vs MEASURED OUTPUT POWER (W)

Figure 47. THD+N vs measured output level.

level reaches the range generally around 0 dBu or 1 V. With a loudspeaker, the level of audible output should be the indication of excessive power levels and with transmitters, the indicated modulation level. If the device under test is not showing any signs of stress, the input level can be further increased in fixed steps of 6 or 10 dB to see if the output increases by that exact amount at each step. When the output fails to increase by the exact amount of input level change, you have exceeded the linear operating range of the device. Both the input and output levels corresponding to the maximum linear level should be noted. Devices with little output power capability such as equalizers, modulation processors, line amplifiers, preamplifiers, mixing consoles, and digital and analog tape recorders will seldom be damaged even at input levels which drive them into severe clipping, so these precautions are typically unnecessary if the device is known to fall into one of these categories.

When a safe output level has been arrived at, the gain of the device should be noted by comparing output and input levels. Some analyzers measure voltage gain directly by providing a real-time ratio mode and internal cable connections which let one analyzer channel measure the generator output while the other measures the device output. A frequency response measurement may then be carried out by sweeping generator frequency at the fixed output level just determined. If the response of the device causes the output to rise significantly at any frequency compared to the maximum value of linear operating range already determined, the generator level should be reduced and another response sweep made. This process may need to be repeated several times until a response sweep is made which does not exceed the proven linear output range

of the device, in order to know that the device was not clipping at frequency response peaks.

THD+N versus frequency measurements are valuable indicators of performance across the audio spectrum. If the distortion curve is essentially constant at all frequencies, this is often an indication of noise-limited performance. Most noise-limited devices produce their lowest distortion just below the clipping point, if this has previously been determined to be a safe operating level.

Performance across the dynamic (amplitude) range of the device can be determined by amplitude sweeps while measuring THD+N and output level. Sophisticated two-channel audio test sets can measure both these parameters on a single amplitude sweep. The upper end of the generator amplitude range should not exceed the already-proven maximum linear range unless you are certain that the device will not be harmed by higher level signals. Figure 46 shows the results of such an amplitude sweep, with THD+N and level separately plotted as two curves versus generator output level plotted on the horizontal axis. The output power curve inflection at the clipping level is characteristic of a simple linear device. A compressor is more likely to shift to a different (but not horizontal) slope with no accompanying high distortion values above the knee of compression. Figure 47 shows another way to portray the same data, with THD+N now plotted vertically against power output on the horizontal axis with generator output level not shown. This second data presentation is typically the most useful with devices such as power amplifiers.

Input impedance can be estimated by noting the change in 600 Ω generator output level while connecting the generator. The graph of Figure 42 may be used to estimate input impedance. Output impedance can be estimated from the change in output level when switching between the high-impedance (typically 100,000 Ω) and 600 Ω terminated input conditions of the analyzer. Figure 44 can be used to estimate the output impedance. Noise can be measured by removing the input signal and noting the resulting output level. A bandwidth of 22 Hz–22 kHz (or 20 Hz–20 kHz) is appropriate for the noise measurement in most cases.

Figure 48. Balanced input circuit.

Interface to the Device Under Test

　　Connections between the Device Under Test (DUT) and the audio test instruments are important. Improper connections can result in the introduction of noise sufficient to produce inaccurate or useless results.

Balanced Devices

　　In professional audio and broadcasting applications, most devices under test have balanced (symmetrical to ground) input and output circuitry; see Figure 48. There will typically be three terminals or connector pins, sometimes marked with a plus, minus, and ground. The only significance of the plus and minus, if so marked, is that the device output should be in phase with the signal at the "plus" input terminal. The audio signal is applied differentially across those two terminals. The advantage of balanced circuitry is that any noise coupled into a cable between devices or inside a device tends to be coupled into both conductors or terminals with about the same voltage and in phase. The balanced-input stage following will tend to cancel out the in-phase (common mode) noise signal, while properly processing the balanced (differential) signal. The ability of a balanced input to reject in-phase noise is called its common-mode rejection ratio (CMRR). Values of CMRR of 60 dB or more (1,000:1 voltage ratio) are readily available. A 60 dB CMRR means that 1 V of common mode noise voltage will be reduced to the same effect as 1 mV of differential audio signal.

　　Balanced inputs and outputs can be achieved either by electronic methods or by coupling transformers. Electronically-balanced circuits can have excellent balance to ground as long as the magnitude of the noise between either terminal and ground does not exceed a value called the "common mode voltage."

Figure 49. Shielded, twisted-pair cable.

This common mode voltage is typically a small-to-moderate fraction of the dc power supply voltage, often on the order of a few volts. If the actual noise volt-age to ground exceeds this value, the electronic stage will be carried into clip-ping (saturation). In certain applications, common mode noise voltages of tens of volts or more may occur. Examples include telephone lines many miles long (especially in proximity to ac mains lines or electric railway lines, etc.). Effec-tive interfacing to such lines requires transformers where the primary and sec-ondary windings are totally isolated from one another, producing a common mode voltage rating equal to the breakdown value of transformer insulation.

To properly connect test equipment to balanced devices, the test instruments must also have balanced inputs and outputs. Shielded cable should always be used to connect the audio test generator output to the input of the DUT and the analyzer input to the DUT output. The cable used with balanced devices and test instruments should consist of a twisted pair (two conductors) under a sin-gle shield; see Figure 49. This tends to equalize the noise coupled into each conductor as the cable passes through various noise fields. Most commercial cable assemblies will have the shield connected to chassis ground at both ends. This is optimum from a standpoint of rejection of high frequency and RF inter-ference. Theoretically, power mains-related hum problems due to ground loops will be minimized if the cable shield does not connect to both the DUT and the test equipment. However, with balanced devices and test equipment, ground loops should not be a problem. In case of severe problems, breaking all cable shield connections between DUT and test instrument and then making a sepa-rate large-conductor ground connection between the chassis of the DUT and the audio test set may be optimum.

Unbalanced Devices

Consumer audio equipment typically operates with unbalanced input and output circuits. The input or output terminal is "hot," working against chassis ground as its reference. Unbalanced connections between devices are normally acceptable only with short cable runs (1 meter/3 feet or less) in relatively noise-free environments, or in very low-impedance high-amplitude applications such as the cable between a power amplifier output and a loudspeaker.
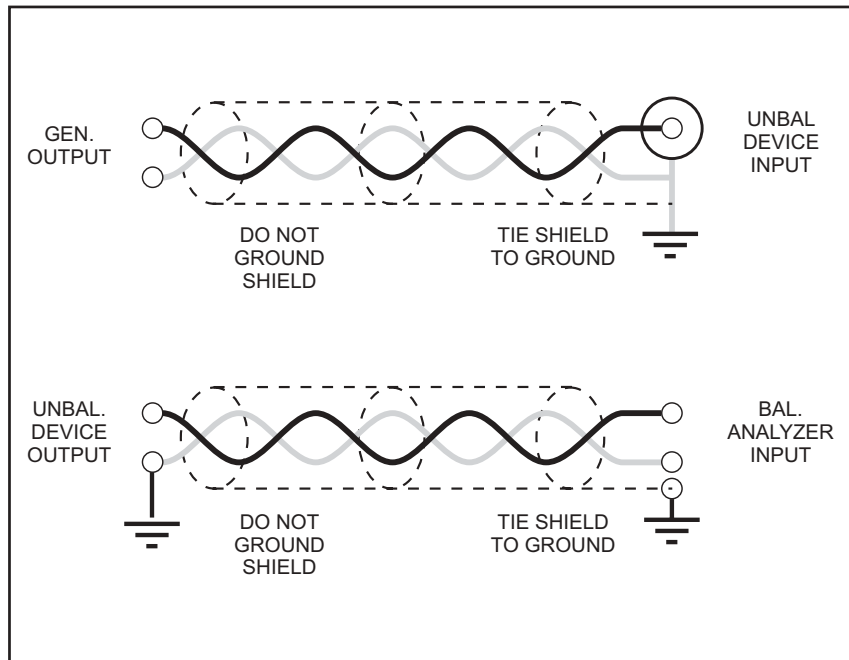
Figure 50. Connections to unbalanced device under test.

Cable connections between unbalanced devices and audio test equipment should be made with shielded twisted pair (two conductors under shield). If the test generator has an unbalanced output configuration, that should be selected. High-quality audio analyzer inputs are always balanced, even when measuring at the output of an unbalanced DUT. To avoid creation of ground loops with their associated hum problems, cable shields should connect to chassis ground only at one end. The cable from test generator to device input should have its shield connected to chassis ground only at the DUT input end. The cable from DUT output to audio analyzer should have its shield connected to chassis ground only at the audio analyzer input end; see Figure 50. The rationale is that DUT and analyzer inputs are typically the most sensitive and highest impedance points, where a direct chassis ground minimizes problems due to stray capacitances, etc. When a stereo device is being tested, the cables from left and right channels should be dressed tightly together rather than separated,

to minimize the area of the loop they form and thus minimize magnetic coupling of noise into that loop. In difficult cases, a separate chassis bonding connection should be tried between DUT and test instrument.

Digitally-Implemented Advanced Measurement Techniques

Most of the audio test techniques discussed up to this point have been in use for several decades. These techniques are generally realized by hardware, largely analog hardware. Recent advances in digital techniques, especially the Digital Signal Processor (DSP), have made more advanced techniques available.

When an analog signal is to be measured by a digital technique, it must first be converted into the digital domain. Such analog-to-digital conversions have been widely described and will not be discussed in detail here. Briefly, the conversion process involves passing the signal through a low-pass (anti-alias) filter, sample-and-hold amplifier, and analog-to-digital (A/D) converter; see Figure 51. The anti-alias filter prevents passage of signals above one-half the sample clock rate, since such signals will produce alias responses which cannot later be removed. The converter output is a series of binary numbers at the sampling rate, with the value of the binary number corresponding to the instantaneous voltage of the audio signal at the instant of sampling.

Digital domain signals can be mathematically processed in any computer. A DSP is a specialized computer optimized for the types of computations typically involved in signal processing and analysis. All DSP operations can be broadly grouped into two classes—real time processes, and "batch mode" processes. A real time process is one which can be carried out in the DSP at the in-
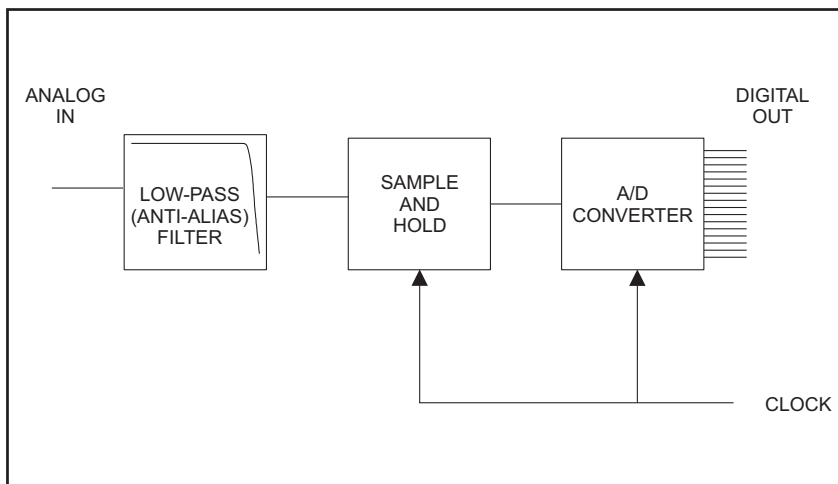

Figure 51. Block diagram, analog-to-digital conversion.

coming sample rate. For example, a DSP may digitally filter an audio signal by accepting an input sample, computing the output result, and passing it on to the following device in time to accept the next input sample. Batch mode processes are those processes in which a number of samples must first be acquired into a block of memory before the process can begin. The DSP then performs some computations (such as a Fast Fourier Transform—FFT) on the memory contents, and finally the results of the computation may be used.

Examples of real-time measurement techniques using digital or DSP technology are principally filtering and level measurement. A powerful DSP unit may compute quite sophisticated filters at input sample rates sufficient to handle full audio bandwidth signals. For example, test equipment is available which can compute ten-pole filter responses in real time at a 48 kHz sample rate. The filters can be high pass, low pass, bandpass, or weighting filters.

The principal measurement examples of batch-mode DSP processes are waveform display and spectrum analysis by FFT. Both start by acquiring a sequence of digitized samples of the signal into memory; this set of samples is called the record.

Waveform Display (Digital Storage Oscilloscope Mode)

Waveform display is relatively simple in concept. Samples from memory are sent to the display device and displayed such that time (sample position relative to the starting point of the record) increases to the right along the horizontal axis, while the instantaneous signal amplitude is plotted at the corresponding points on the vertical axis as in any oscilloscope display. Samples from the record may require mathematical conversions and possibly digital-to-analog conversions to the format needed for display. If the display device is a cathode ray tube, the digitally-formatted samples must be converted by a D/A converter to the analog voltages required to deflect the electron beam. If the display device is a computer screen, the processing will only involve changing from the memory storage format to the digital format required by the display.

When the portion of the record to be displayed consists of hundreds of samples or more, it may be satisfactory to directly display each sample as a point. In other cases, additional digital processing is generally desirable for waveform display, rather than simply sending the stored samples to the display. Often the portion to be displayed contains only a small number of samples, as in the case of viewing a few cycles of a waveform whose frequency approaches one-half the sample rate. Simply plotting those samples as points and connecting them with straight-line vectors produces a very distorted version of the waveform which may not even be recognizable (see Figure 52). In this case, it is desirable for the DSP to interpolate additional values between the actual data record samples (see Figure 53). Interpolation algorithms are based on the fact that the record consists of samples of a band-limited signal.

AUDIO PRECISION 17 kHz SINE ACQUIRED AT 48 kHz SAMPLE RATE, NO INTERPOLATION
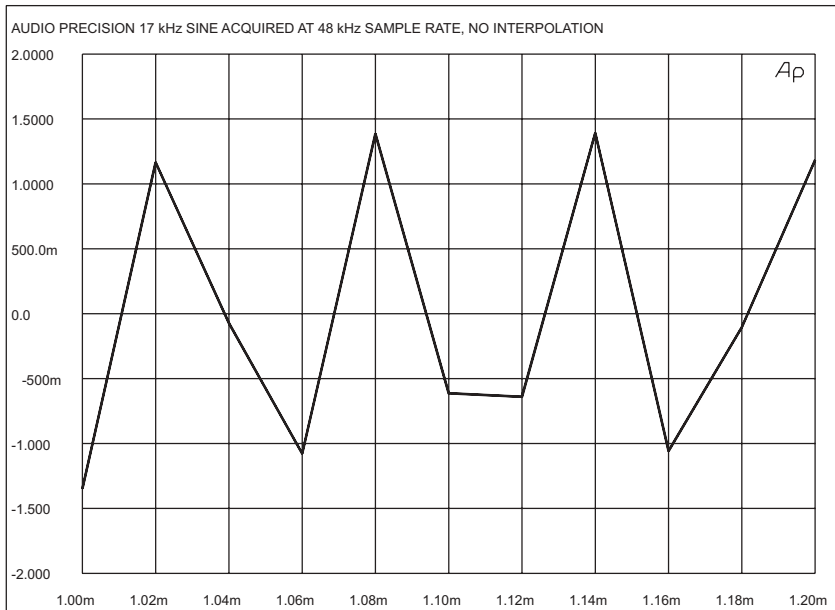
Figure 52. Time domain waveform display, signal frequency approaching Nyquist frequency; no interpolation.

AUDIO PRECISION 17 kHz SINE ACQUIRED AT 48 kHz SAMPLE RATE, INTERPOLATED

Figure 53. Same waveform as Figure 52, interpolated.

Figure 54. Envelope of burst signal.

Other types of processing can be valuable when it is desired to view a wide dynamic range, which is often the case in audio. A linear vertical display, standard with normal oscilloscopes, makes it very difficult to distinguish much detail below approximately 1/30 to 1/100 of the maximum peak-to-peak amplitude of the signal (–30 to –40 dB). If the signal is processed to produce the absolute value (always a positive value) of the peak voltage over each half cycle, then logarithms can be used and the signal envelope can be displayed using decibel units. Figure 54 shows an audio burst signal in both linear and dB vertical modes, with the dB display giving useful information down to levels much farther below signal peak amplitude.

Fast Fourier Transform (FFT)

The FFT is an efficient method of converting a digital amplitude-vs-time record into an amplitude-vs-frequency spectral display. There are several basic relationships necessary to understand in order to use FFT-based instruments effectively:

- as in any sampled process, the maximum upper frequency limit is one-half the sample rate

- the FFT process will produce an amplitude-vs-frequency data set with half as many points ("bins") as the number of samples in the amplitude-vs-time record from which it was made

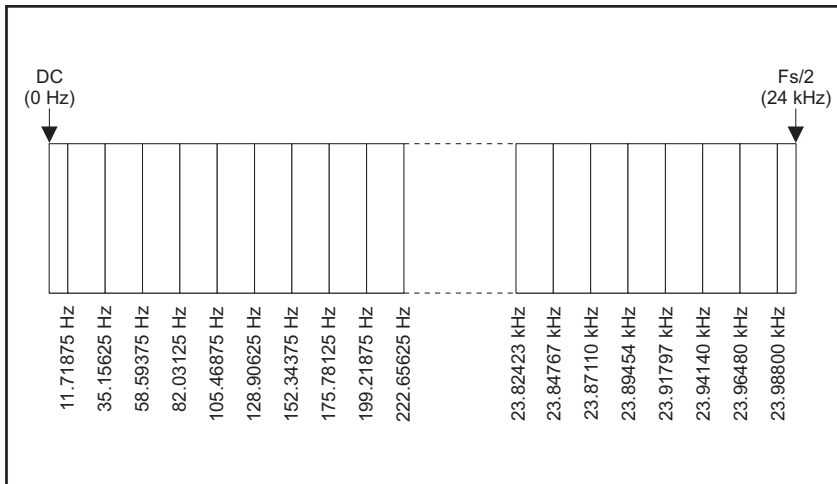*Figure 55. Example, FFT spectrum analysis bins with 48 kHz sample rate, 2048 waveform samples, thus 1024 FFT bins.*

■ an FFT is a frequency-linear process; all bins are the same width, extending from 0 Hz (dc) to ½ the sample rate.

For example, assume a 2048-sample amplitude-vs-time record acquired at a 48 kHz sample rate. Performing an FFT on this record will produce 1024 amplitude values, each representing the amplitude of the spectral products within a frequency bin. One half the sample rate is known as the Nyquist frequency. With a Nyquist frequency of 24 kHz and 1024 linear bins, the width of each bin is 24,000/1024 or 23.4375 Hz. The first bin is centered at DC (zero Hz) and extends from –11.71875 Hz to +11.71875 Hz. The amplitude value for the second bin represents the amplitude of signals between 11.719 Hz and 35.156 Hz, etc.; see Figure 55. Better frequency resolution can be obtained by transforming a longer record or acquiring signal at a lower sample rate, but the lower sample rate will also reduce the maximum signal frequency which can be acquired.

The FFT process inherently assumes that the portion of signal acquired as the record is an exact sub-section of an infinitely repeating signal. Stated another way, the FFT process assumes that if the end of the record were "spliced" to the beginning of the record and this record used to generate a continuous signal, the signal generated would be exactly like the signal originally sampled and acquired. In practice, unless the signal to be acquired is deliberately generated in some synchronous fashion with the acquisition process, this will not be the case. When a random or independent signal is acquired, it is extremely unlikely that all frequency components of the signal will be at the exact phase and amplitude at the end of the record necessary to make a transient-free splice to the beginning of the record. Performing an FFT on such a non-synchronous signal essentially gives the spectrum as if the signal were pulsed, with the energy apparently being spread across the frequency spectrum rather

than concentrated at the specific signal frequencies; see Figure 56. A solution to this problem of apparent spectrum spreading of non-synchronous signals is the "window."

A window is a magnitude function of time. Most windows have an amplitude of unity (1.000) at their center and an amplitude of zero at the beginning and end. The first window used in FFT work was proposed by Hann and is a raised cosine, as shown in Figure 57. When each amplitude sample value of the record is multiplied by the window magnitude at each corresponding time point, the result is a new record with its center values essentially unchanged while the values at beginning and end are forced to zero; see Figure 58. Now, the end and beginning of the window-multiplied record "splice" together smoothly since they are both zero. The FFT of this window-multiplied record then shows no energy spread across the full spectrum. However, the window-multiplied record clearly is not the same as the original signal. This change is evidenced in the resulting FFT by spillage into adjacent bins; a pure sine wave which should be represented by energy only in one bin instead shows energy across a number of adjacent bins. A variety of different window shapes have been proposed, each of which provides different compromises of possible amplitude measurement error when the signal frequency is not exactly at bin center, and different amounts of energy spillage into adjacent bins. Figure 59 shows the selectivity of several common windows. The Hann window has good ability to separate signals of nearly equal amplitude which are quite close together in frequency, but its selectivity further away from center frequency is


*Figure 56. FFT spectrum analysis of nonsynchronous sinewave, no window.*

Figure 57. Hann window.

Figure 58. Continuous sine wave signal multiplied by Hann window.

not as good as others. The Blackman-Harris window has excellent selectivity a few bins away, although not quite as sharp as the Hann window at the very center. The Flattop window has poorer selectivity than either Hann or Blackman-Harris, but has essentially zero error when measuring the amplitude of signals falling anywhere in the bin. Conversely, the Hann window can display an error of as much as 1.5 dB and the Blackman-Harris window an error of 0.8 dB if the signal frequency should fall at the extreme edge of the bin rather than at exact bin center.

A useful capability in digitally-based analyzers is to be able to acquire a relatively long record of test signal or program material, and then to perform FFT analyses at different points in the acquired record. When both the starting point and the length of record to be transformed are selectable, frequency analysis of complex signals can be made very flexibly. For example, if the test signal is a sine wave burst from a low level to a higher level and back to the low level, the output of compressors and other dynamic processors can be analyzed to measure distortion during the attack, compression, and release times of the compressor.

FASTTEST Technique

A combination of a special digitally-generated signal and special FFT-based analysis technique have been recently developed called FASTTEST. The signal consists of multiple simultaneous sine waves (multi-tone). A typical signal

Figure 59. Response of BH4, Hann, Flat and No windows. Horizontal axis calibration is delta-Hz from signal frequency; bin width is 2.93 Hz.

Figure 60. High-resolution FFT of FASTTEST multitone signal.

might consist of about 15 sine waves spaced evenly across a logarithmic graph of the audio spectrum, although larger and smaller numbers of tones and different spacings are possible. At the output of the DUT, this signal is acquired and an FFT performed; see Figure 60. If the multi-tone signal is digitally generated from a generator waveform buffer exactly equal to or exactly one-half (or one-third, one-fourth, etc.) the length of the record acquired into the analyzer, and if all of the sine waves in the multi-tone signal are exactly synchronous with the generator buffer, no window function is needed in the analyzer. In this case, the effective measurement selectivity is the analyzer bin width and no energy spillage occurs into adjacent bins.

After acquisition and FFT, the resulting set of amplitude bins can be further analyzed in a variety of fashions for different purposes. If we examine only the analysis bins into which the fundamental frequencies fall, we have obtained the frequency response of the DUT. If we integrate the amplitudes of the bins between the fundamental frequency bins but not including the fundamental bins, we have a measurement of total distortion (harmonic plus intermodulation) and noise. If we generated the signal from a generator buffer exactly one-half as long as the analyzer record length, the analyzer frequency resolution will be twice the generator resolution. All generator fundamentals and all harmonic and intermodulation products of those generator fundamentals must fall at frequencies which are exact integer multiples of the lowest frequency sine wave which could be generated from this buffer. The analyzer has twice as many bins, however, with alternate bins falling in between the frequencies where generator-related signals can occur. Measurement of the amplitudes of these "empty" bins produces a measurement of noise in the presence of test signal.

If the DUT is a stereo device, the FASTTEST technique can also measure inter-channel phase. FFT analysis produces the phase of each signal. If the phase of one channel is subtracted from the phase of the alternate channel, the result is inter-channel phase. It is also possible to measure stereo separation (crosstalk) with the FASTTEST method. Two different waveforms, each containing some frequencies not found in the other, are fed to the two channels of the stereo device. The frequency bins of each channel corresponding to signals fed only to the opposite channel may then be examined to see how much signal has "leaked" across the channels.

Figure 61. FASTTEST Signal shaped to emulate spectral energy distribution of program material.

A key advantage of the FASTTEST technique is speed. As little as 1–2 seconds of multitone signal can be fed to the device or system under test and captured into the analyzer. Where absolute minimum test signal duration is required, such as testing broadcast links during normal programming periods, the FASTTRIG variation on FASTTEST can recognize and capture a multitone burst of as little as 270 milliseconds. The device or system under test may then be restored to its normal uses while the FFT is made and the FFT data is examined for various parameters such as frequency response, distortion, noise, phase, and separation. Another advantage of FASTTEST is that the multi-tone stimulus signal is much more similar to music and voice than is a single sine wave. In fact, the signal can be made even more similar to program material by making each tone the appropriate amplitude so that the spectral energy distribution is identical to the distribution of the desired type of program material; see Figure 61 for an example. This type of spectral shaping is particularly valuable when measuring performance of sophisticated multi-band signal processors such as the modulation level enhancement devices used by most broadcasting stations.

Digital Domain Testing

Recent years have seen a large shift to digital techniques in the implementation of various audio processes. The digital compact disc has nearly totally replaced the analog vinyl long-playing phonograph record and has also taken significant market share from analog tape cassettes as a medium for pre-recorded music. Digital recorders of several varieties, including RDAT (Rotating head Digital Audio Tape), are widely used in broadcasting and recording studios for mastering. Digital direct-to-computer-disk recorders and editors are used instead of analog tape cut-and-splice techniques for most editing and preparation of sophisticated sound for film and television. Artificial reverberation and echoes are now produced almost exclusively by digital techniques rather than the springs and plates used in the past. All of these changes and more have brought the need for testing in the digital domain or mixed analog and digital domains.

To intelligently design and test digital audio equipment and to trouble-shoot a failure or performance degradation to a replaceable block requires the ability to generate both analog and digital domain signals and to make measurements in both analog and digital domains. Input and output amplifiers are analog and may be tested by traditional techniques. An A/D converter must be tested with an analog input signal and digital domain analysis techniques. A D/A converter requires a digital domain stimulus signal and analog analysis. Measuring performance of the central digital signal processing or recording mechanism should be done purely in the digital domain.

Digital Audio Formats

Linear pulse-code-modulated (PCM) digital audio signals may be transmitted, stored, and processed in a wide variety of formats. Since parallel transmission of a stereo 16-bit audio signal requires a minimum of 32 conductors plus grounds, parallel interconnect is rarely used except for possible short runs between chips on a circuit board inside a piece of equipment. The digital audio interface of early Sony digital processors and recorders was the proprietary Sony SDIF and SDIF2 serial format, which used three coaxial cables.

At present, digital audio signals are commonly sent between professional and broadcast devices in the AES/EBU format. AES is the Audio Engineering Society; EBU is the European Broadcasting Union. In consumer equipment, the very similar SPDIF/EIAJ format is used for inter-equipment interface. SPDIF stands for Sony Philips Digital Interface; EIAJ is the Electronic Industries Association of Japan. Both AES/EBU and SPDIF/EIAJ are serial formats using a two-conductor cable. In the AES/EBU format, the two conductors are a balanced twisted pair under a shield and terminated in "XLR" or "Cannon" connectors, as typically used for microphone and audio line connections in studios. In the consumer format, the two conductors are typically an unbalanced
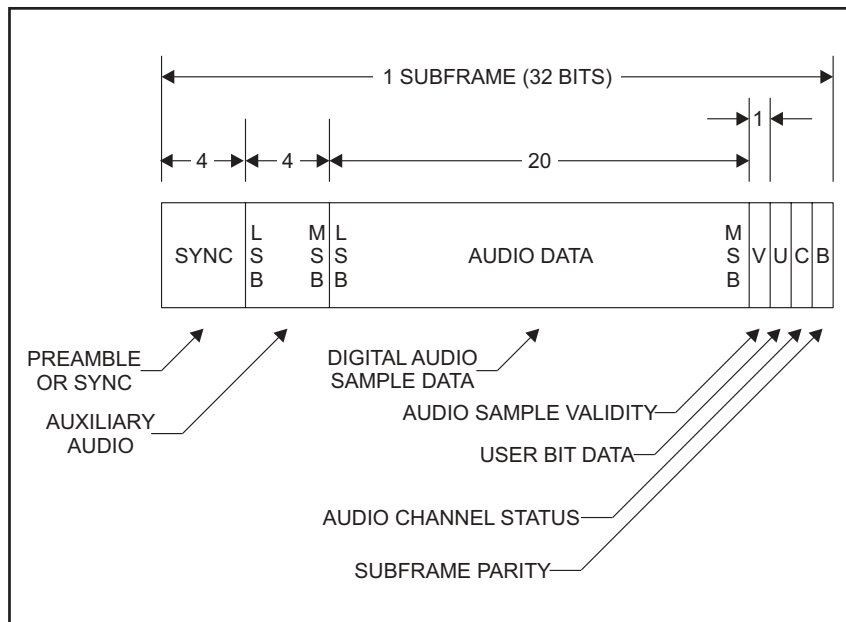
Figure 62. AES/EBU subframe structure.

single coaxial cable terminated in "RCA phono" or "Cinch" connectors, or a fiber optic cable. In both cases, all the bits making up one stereo pair of audio samples plus ancillary data are serially transmitted down the cable. The complete set of bits representing a left channel sample, right channel sample, and other data and synchronizing information are called a "frame," and frames are thus transmitted at the sampling rate of the program material.

The three standard sampling (frame) rates are 48 kHz, 44.1 kHz, and 32 kHz. Each frame consists of two subframes, one for channel A (left) and one for channel B (right). Each subframe consists of 24 bits for audio (or auxiliary data), plus 8 bits for preamble, synchronization, audio sample validity, user bits, channel status bits, and subframe parity. Present consumer systems such as Compact Disc and RDAT only use 16 of the possible 24 bits for audio, but the standard allows for up to 24 for improved systems.

The channel status bits are assembled over blocks of 192 frames into 24 bytes called Channel Status Bytes. Figure 62 is a representation of an AES/EBU sub-frame. Many items of equipment with digital input require a specific value for the channel status bits before they will function properly. Digital recorders, for example, can typically not be made to record unless the proper channel status bits are sent as part of the signal. It is often necessary for the channel status bits to agree with the actual sample rate according to the AES/EBU or SPDIF/EIAJ standards. Thus, audio test generators with digital outputs typically also require the ability to set user-selectable values into the channel status bits.

These inter-equipment interconnect standards are not commonly used within an item of equipment for storage or processing. RDAT, stationary head digital recorders, and Compact Discs each have their own unique methods of storing and reproducing the digital data. A/D and D/A converters and DSP chips use many different techniques for the representation of digital audio data. Most chips of 16 bits and higher resolution use some form of serial interface, with the Philips I²S being used by several other manufacturers. There is no effective industry standardization at the chip level, so many companies and chip designers have chosen unique formats convenient for their own purposes.

Large amounts of digital storage or bandwidth are necessary for the recording or transmission of linear PCM at the 16-bit (or higher) resolution level and the 44.1 kHz or 48 kHz sample rates necessary for full 20 kHz audio bandwidth reproduction. For example, a 16-bit 48 kHz sample rate stereo signal with no ancillary bits for synchronization or other purposes amounts to 1,536,000 bits (192,000 bytes) per second. Transmission of this full signal would require wide bandwidths. Recording one hour of such a signal requires about 690 megabytes of data storage. Much work has been done in the latter half of the 1980s and early 1990s on reducing these bandwidth or storage requirements with minimal compromise of audio quality in order to make digital audio broadcasts practical or provide more economical recordings. The Philips Digital Compact Cassette (DCC) and Sony MiniDisc (MD) are the two best-known consumer products in this area. NICAM, MUSICAM, Dolby AC-2, APT-X, and a number of other techniques are well-known among broadcasters and audio professionals. Generically, these systems and devices are often referred to as "codecs," "digitally-compressed audio" or "low bit-rate audio." Processing and storage or transmission are intimately related in each of these systems, so each has its own format.

Generation of Digital Domain Signals

Signal generation in the digital domain can be done by dedicated digital hardware, but is most flexibly accomplished by use of DSP (Digital Signal Processor) technology. Signals can be generated by several methods, including real-time computation of waveform samples, storing one-fourth of the waveform in memory and manipulating phase as it is generated, or via an arbitrary waveform architecture where an entire waveform is stored in a memory buffer and clocked out one sample at a time. Non-DSP digital generators are typically limited in their capabilities. They often can generate only sine wave signals and only at a few specific frequencies related to the sample rate, and may have relatively poor amplitude resolution. DSP-implemented generators can have very high frequency resolution (better than 0.01%) and amplitude resolution at the 24-bit level, and can generate complex signals.

Analysis of Digital Domain Signals

Analysis of digital-domain signals can be broadly grouped into two categories—real time analysis and batch-mode analysis. In real time analysis, each digital word at the incoming sample rate is processed through a series of digital filtering and digital amplitude measurement functions to produce an output value for display or graphing. Real time digital domain analysis can have the same "look and feel" as analog analysis and appears to the operator no different than conventional hardware analysis of analog domain signals. Swept measurements can be made and graphed, or some form of real-time "meter" indications (usually on a computer display screen) can be observed while making adjustments to the digital device under test. Performance of real-time digital measurements is primarily limited by the power of the DSP chip, since only the finite time between input samples (20.8 µs at a 48 kHz sample rate, for example) is available to compute the filtering and level measurement functions. With DSP chips such as the Motorola 56001, this is sufficient time to implement ten-pole filter computations which can produce filters comparable or superior in performance to the analog hardware filters used in better analog audio analyzers.

Batch-mode analysis of digital domain audio signals consists of several sequential steps: acquisition of a block of consecutive digital samples into memory, optionally performing a computation (typically an FFT) on that block, and then displaying or otherwise using the results of the computation. For simple waveform display, the value of the samples may be graphed versus sample position (time). This mode becomes a digital domain digital storage oscilloscope. If an FFT is performed and displayed, we have a digital domain spectrum analyzer. The FASTTEST technique uses a digitally-generated complex waveform, FFT analysis of the results, and additional post-processing. Thus, FASTTEST can be used with any combination of digital domain and analog domain inputs and outputs.

Specific testing techniques using these digital domain tools plus more conventional analog domain techniques will be discussed under digital applications in Section 3.

Section 2

Major Testing Environments

Most testing takes place in one or another of three environments or activities: research and development in the laboratory, production test and quality assurance on the factory floor, and maintenance and repair in the field or in a service shop. Other activities can be listed such as on-site acoustical measurements in order to optimize an auditorium or theater, but the three major areas stated possibly cover 95% of all audio testing. Each of these activity areas tends to emphasize certain aspects of testing and minimize others, as will be summarized in this section.

Laboratory Research and Development

Laboratory testing tends to be the most comprehensive of the three major environments. A development engineer must measure almost every conceivable parameter of a new design in order to verify that it will meet the development project goals and assure that there are no surprises. Thus, a development engineer is likely to test a product for THD+N, two or three forms of intermodulation distortion, and also perform a number of FFT spectrum analyses at different stimulus frequencies and amplitudes in order to thoroughly explore nonlinearities of the device. He is likely to measure THD+N as a frequency sweep at several amplitudes including maximum rated amplitude, and also as amplitude sweeps over the full dynamic range of the product, possibly at several different frequencies.

Development testing results are likely to be displayed graphically whenever possible. Graphing test results makes it simpler to perceive relationships between the data points and thus to interpret the data. Permanent documentation, usually printed, is also a typical requirement of development laboratory testing. This documentation can help compare performance of several design alternatives and will form the basis for production test specifications when the device moves into production.

Development testing consists of both structured and unstructured phases. When a prototype fails to initially deliver the expected performance, the engineer is likely to be led by the results of one measurement into the next type of measurement to be made. A hint of an irregularity in response or distortion at a certain frequency may lead to high-resolution exploration of performance near that frequency, but over a wide dynamic range or with different types of measurements. Other development testing examples may be very structured. These can occur early in a project when a number of samples of a new component may be subjected to exhaustive testing over a range of temperature and humidity, or late in a project when a completed prototype is run through a completely-structured battery of tests.

Development testing tends to require the highest performance test equipment available. The development engineer wants to know the actual measurements of his device, not merely whether it is better or worse than some acceptable limit value. The development engineer also does not know what the next project may be and therefore what level of performance may be required in his test equipment for that next application.

Production Test and Quality Assurance

This area is the most structured of all testing areas. By the production phase, a test specification exists with rigid numbers which must be met for acceptable performance. Due to the economic impracticality of testing everything that was tested in development, production testing normally covers only the most important performance parameters of a device. Performance parameters are also likely to be measured at fewer points; where the development engineer makes a 100-point sweep of frequency response, the production test procedure may only measure at 20 Hz, 1 kHz, and 20 kHz. Some test philosophies include abbreviated testing on all units built, plus more exhaustive testing or audits on randomly selected samples.

Production testing is the easiest and most worthwhile type of testing to automate. The greater testing speeds available through automation allow management to have better confidence in the quality of their products, by testing more parameters or testing each parameter in more detail. Test automation can produce necessary documentation at virtually no cost and easily supplies results necessary for statistical evaluation of product quality. Except for very low production volumes, most companies today have automated their production tests. Automation requires test instruments with computer interfaces plus appropriate software running on the computer. A number of manufacturers of computer-controllable audio test instruments also supply software packages, either included as standard or at an extra cost. Some of these software packages require no prior experience in programming in order to set up automated test routines. To be useful for production test work, these software packages require a number of features:

- The ability to create swept tests across a specified frequency range or amplitude range, with the test points either computed automatically or looked up from a user-specified table.

- The ability to specify upper and/or lower limits of acceptable performance for each test and each device.

- The ability to have the computer make go/no-go decisions based on whether test data falls within the specified limits.

- The ability to link together a number of different tests such as response, distortion vs frequency, noise, etc., into a complete test procedure.

- The ability to put messages onto the computer screen to prompt the operator into actions which must be performed manually.

- The ability to send commands to some general-purpose interface in order to control modes and functions of the device under test, other instruments, or other conditions and actions of other portions of the test setup.

- The ability to branch to other parts of the test routine, based on a device passing or failing a test.

- Support for printed documentation in a variety of formats (graphic, tabular alphanumeric, etc.) to a variety of output devices such as dot matrix and laser printers.

- The ability to save test results to disk.

Depending upon the test volume and acceptable test times, an automated factory test system may need to connect to device handlers or robots which can place the DUT into a test fixture. In cases of complete device functional text, test fixtures may only need to connect to the DUT input and output. In other cases, it may be necessary for test and adjustment purposes to access internal circuit nodes and a "bed of nails" test fixture may therefore be necessary.

It is also common in more sophisticated cases of test automation for the individual computers which control each test station to be connected by a network into a central host computer which will keep track of test data in order to compute statistics of the process. The central host sometimes also permits downloading of new tests or new acceptance limits.

Maintenance and Repair

Testing activities in maintenance and repair may take place in service and maintenance shops and depots or in the field, at the actual usage location of the device under test. These activities may include alignment and adjustment for optimum performance, complete tests to verify performance (often called "proof of performance" tests), and the trouble-shooting necessary for fault location to the replaceable device, module, or component level.

Thus, maintenance and repair testing may vary from very structured to totally unstructured. Proof of performance tests are typically as structured as factory production tests, going through specific parameters under specific conditions and comparing the results to specific limits of acceptable performance. Troubleshooting tends to be at the opposite extreme of structure, where the experienced service engineer or technician makes informed guesses as to causes of the symptoms and is steered by the results of one test to the next logical test to make in order to narrow down the source of the fault.

Diagnostic techniques are particularly important in this type of testing. Spectrum analysis—either low resolution via swept bandpass filter or high resolution via FFT, helps locate noise and interference sources such as ac mains-related hum or signal coupled into an audio device from a video or computer monitor. Graphic display of results is helpful to quickly point out problem areas such as clipping at a lower level than normal or a frequency response problem.

Section 3

Testing Specific DUT Types

Introduction

The following sections will deal with specific types or families of audio devices under test. The tools and techniques described in the first device category (amplifiers, etc.) will be applied in many of the following device categories. Specific variations and specific emphasis on particular tools and techniques will be discussed. The range of expected performance measurements for good and excellent performance on each type of device will also be discussed.

For virtually all types of audio devices to be tested, a fairly standard battery of tests is normally performed. This standard test battery includes frequency response, distortion, noise, maximum output capability, and (if a stereo device) separation or crosstalk. Certain more specialized measurements are typically added to that standard test battery when the class of device under test is known to have certain potential weaknesses or limitations.

Amplifiers and Similar Real-Time Linear Analog Audio Devices

This category of audio device is the most common. It includes power amplifiers, pre-amplifiers, microphone and phono amplifiers, line amplifiers, distribution amplifiers, equalizers, most portions of mixing consoles, etc. All these devices have in common the fact that they are nominally linear and produce a real-time output similar to their real-time input, but typically at a higher power or voltage level or a different impedance environment and often with a modified frequency response. This category also provides the basic testing method discussion for all the DUT categories discussed later. Companders, analog tape machines, radios and tuners, etc. are tested in the same manner as amplifiers except for the specific variations, considerations, and additional tests discussed in each of those sections.

Standard Tests—Response and Distortion

The standard battery of tests is straight-forward for this category of devices. The test oscillator drives the DUT input while measurements are made at the output. Generator output amplitude is normally held constant during frequency sweeps. Time delay through the device is normally negligible. Frequency response measurements are typically made near the maximum output level capability of the device. If the DUT has large variations in frequency response across the frequency band, care must be taken to not test at such a high level that the output will be driven into clipping or saturation at the frequencies with highest gain. Amplifiers can easily be designed whose response varies less than ±1 dB across the 20 Hz–20 kHz audio spectrum. Flatness of ±0.1 dB is achievable in excellent designs. Amplifiers which are likely to be cascaded, such as distribution amplifiers in a broadcasting studio or recording facility, may have even tighter response such as ±0.05 dB.

Measurements of harmonic distortion (THD+N) versus frequency are typically made across the entire audible spectrum, particularly if the device is not sharply band-limited at the upper frequency limit. Low-level devices such as equalizers and line amplifiers normally exhibit lower values of THD+N than power amplifiers. Good low-level amplifier designs routinely produce less than 0.01% (–80 dB) distortion across the entire 20 Hz–20 kHz spectrum at levels 1 or 2 dB below their clipping point. If the devices have transformer output circuitry rather than electronically-coupled, the distortion figures are likely to be somewhat higher and in particular to rise at low frequencies and high power levels as the transformer approaches core saturation. Thus, a transformer-output device with 0.01% distortion from 1 kHz to 20 kHz may rise to 1% or more at 20 Hz.

Graphs of output THD+N versus an amplitude-swept input signal at a fixed frequency (as shown earlier in Figure 47) are also quite useful. Such graphs show the point at which the device goes into clipping. Many modern, well-designed devices in this category are noise-limited over most of the dynamic range below clipping. Any distortion generated in the device is usually below the wideband noise level until virtually the onset of clipping. Thus, the THD+N versus amplitude curve in percent or dB units will show the characteristic rise with decreasing level, since the "N" component of THD+N dominates. As previously discussed, this does not indicate that noise is increasing at lower test signal amplitudes; the noise level is typically constant and independent of amplitude. When THD+N is expressed either in percentage of the output level or dB below the output level, this constant noise signal represents a larger and larger fraction of the lower amplitude signals. Actual noise and distortion mechanisms are often clearer if the result of such an amplitude-swept measurement is expressed in absolute units—V, dBu or dBV of noise and distortion, as was illustrated in Figure 24 on page 34. This presentation of the same information makes it much clearer that the device has a constant noise

output independent of signal level and, at high amplitudes near and above clipping, actual distortion products rise above the noise floor.

Maximum Output Level

The point of onset of clipping may be defined as the maximum output level capability of the device. Since "onset of clipping" is somewhat ambiguous, maximum output level is often defined as the highest output level achievable without exceeding a certain specified percentage of distortion. Although the term "clipping" comes from observation of a waveform on an oscilloscope, a THD+N measurement is a far more sensitive indicator. When THD+N starts to increase rapidly as input signal level is increased, the waveform clipping is still typically invisible on an oscilloscope.

Noise and S/N Ratio

Noise measurements must specify both the measurement bandwidth (or weighting filter used) and the input termination resistance. Ideally the input termination for noise tests would be equal to the output impedance of the device normally driving this device. In practice, a short circuit (0 Ω) is often specified as the input termination for noise measurements as a matter of convenience. Noise of this class of devices is frequently specified in terms of signal-to-noise ratio (S/N or SNR), rather than an absolute value of noise voltage or power. S/N ratio measurements are made by first establishing a specified output level (usually near the maximum or normal operating output level of the device) and defining this as the 0 dB reference. Signal is then removed and replaced with the specified source resistance, and the resulting change in output level is the S/N ratio. Thus, S/N ratio specifications require a statement of the reference output level in addition to bandwidth and input termination.

Crosstalk and Stereo Separation

The terms crosstalk and stereo separation are often used almost interchangeably in audio measurements. The technique of measurement of crosstalk and of separation is identical; the difference in terminology depends upon whether the two channels being measured are designed to carry completely independent, non-correlated signals (typical of telephone applications) or the strongly-correlated left and right signals of a stereo program. Thus, "crosstalk" should properly be used to describe signal leakage from one channel into another which should be completely independent while "stereo separation" should be used only when the two channels are designed to carry a stereo program. The term "crosstalk" will be used in the generic sense in this document to describe both mechanisms.

Crosstalk should normally be measured as a function of frequency, since the values are likely to vary strongly with frequency. The most common circuit

mechanism causing crosstalk is stray capacitive and inductive coupling. If circuit impedances are approximately constant with frequency, crosstalk caused by a single instance of capacitive coupling will increase with increasing frequency at a 6 dB per octave rate. That is, crosstalk due to single pole capacitive coupling will be 6 dB worse at 2 kHz than at 1 kHz, etc. Crosstalk can also be caused by inductive coupling, shared power supplies, shared ground returns, etc., so the relationship is often not the simple capacitive coupling model.

Crosstalk should normally be measured selectively, with a bandpass filter in the analyzer tuned to the generator frequency, in order to measure crosstalk at or below the wideband noise level. This is not merely an academic concern; the human ear is able to distinguish coherent signals such as sine waves even when the signal amplitude is 10 dB to 20 dB below the wideband noise level. Figure 63 is a graph of crosstalk versus frequency in a two-channel electronic device. One line is crosstalk from A into B; the other line is crosstalk from B into A. Note that the two curves are commonly not identical. Circuit layout and complex stray coupling mechanisms often result in somewhat different values for crosstalk in the two directions.


Figure 63. Example graph of crosstalk (separation) vs frequency, L to R and R to L.

Good values for crosstalk and stereo separation in electronic devices range from 50 dB upwards, with crosstalk in excess of 100 dB achievable by sufficient isolation. When the application is stereo, such values are far higher than really necessary for a full stereo effect. Psychoacoustic experiments have shown that stereo separation of 25 dB to 30 dB at midband audio frequencies

(500 Hz–2 kHz) is sufficient for a good stereo effect, with even less stereo separation acceptable at the frequency extremes. If the two devices are actually carrying independent programs such as two voice circuits, values of crosstalk of 60–70 dB and more are desirable.

Input vs Output Clipping

Overload and clipping can occur at any stage of an amplifier. If the amplifier has no gain control, it is irrelevant to the user which stage clips first since his only method of avoidance is to reduce the input signal level to the device. When the amplifier has one gain control, it is important to know the clipping point of the amplifier before the gain control so that the maximum input level can be kept below this input clipping level. If the input circuitry is driven into overload, the output signal will be distorted regardless of the setting of the gain control. Very sophisticated audio devices such as mixing consoles may have a number of gain controls at various points in a signal path. Intelligent operation of such a device requires an understanding of the various clipping levels.

Measuring input versus output clipping levels of a device such as a mixing console requires a logical process, considerably helped by some prior knowledge of the probable clipping levels at various points of the circuitry. The output stage clipping point can normally be determined by setting the last gain control (master gain) at maximum while driving the input stage at a low level and using no more gain than necessary in the input stages to reach clipping (point of rapid increase in THD+N). The clipping level so determined should be at or above the maximum rated output from the device. The input clipping point is then measured by setting the input gain control to maximum, reducing the output (master) gain until the output level is 10 dB or more below the previously-determined output clipping level, and then adjusting the generator output level for the threshold of clipping as determined by a sharp increase in THD+N.

Equivalent Input Noise

Noise measurements as described in the "tools and techniques" section are measurements of the output noise from a device. For certain applications, it is desirable to know the input noise of a device. This is commonly necessary when designing or selecting microphone preamplifiers or other low-level devices, in order to determine whether the input noise is sufficiently low with respect to microphone output level to provide the desired signal-to-noise ratio. Input noise cannot be directly measured. Instead, it is determined by a computation from the results of two measurements—output noise, and circuit gain. The output noise measurement is then divided by the circuit gain to determine what the equivalent input noise is, making the assumption that all the noise at the output is simply an amplification of the input noise. In a well-designed

preamplifier, this will be the case; noise contribution of any later stages should be negligible compared to the input noise of the first stage.

Transducer Gain vs voltage Gain

Simple voltage gain of a device is the ratio of output voltage to input voltage while operating at a specified signal level within the linear range of the device. voltage gain thus ignores source and load impedances and the power gain of the device. The term "transducer gain" is used to describe gain measurements of a device made while taking into consideration the nominal source and load impedances with which the device is designed to operate. This distinction is particularly important when actual source and load impedances are reactive, or when source or load are deliberately mismatched. Since microphones are normally deliberately mismatched in load impedance, the transducer gain technique may produce gain measurements nearly 6 dB different than the voltage gain technique.

Transducer gain may be measured and computed either by substitution techniques, or directly if the test instrument has the proper units available and is properly calibrated.

Transducer Gain via Substitution

For the substitution technique, the test oscillator output impedance is first set to the value specified as output impedance of the device which normally drives the device under test in regular operation. For example, if the device under test is a professional microphone preamplifier or mic input of a mixing console, it is normally driven by a professional low-impedance microphone with a nominal impedance of 150 Ω. The test generator output impedance is thus selected as 150 Ω. The generator output is then temporarily terminated into a pure resistance equal to its internal impedance—in this case, a 150 Ω resistor. Under these terminated conditions, the generator amplitude is then adjusted to produce a measured level across the termination equal to the specified input level for the test to be made; –60 dBm might be typical for a mic input. The resistive termination is then removed and the generator output connected into the input of the device under test, with no further adjustment of the generator amplitude. Assuming that the output of the device under test is connected into a resistor equal to its rated load impedance, the transducer gain of the device is now computed as the dB difference between terminated output power level in dBm units and the generator dBm output level which had been measured and set while terminated into the matched resistive load. The reason that transducer gain measurements may approach 6 dB difference from voltage gain measurements in this microphone example is that microphone preamplifiers are typically designed with input impedances on the order of ten times higher than the nominal microphone impedance. A 1,500 Ω input impedance is virtually an open circuit to a 150 Ω test generator (or microphone), so the generator

voltage rises to almost twice the value (+6.02 dB) that it was while terminated into a matched resistor.

Transducer Gain via Proper Test Instrument Features

If the test generator has been designed with properly-implemented power units (dBm and/or watts) including the ability for the user to specify to the instrument the value of the load resistance, the complicated process of the substitution technique described above may be avoided. In the preceding microphone preamplifier example, the generator 150 Ω output impedance is selected, the dBm unit selected, a 150 Ω input value specified for the device under test, and the desired level (–60 dBm in the preceding example) supplied as the generator output level. The "smart" generator will set the open-circuit voltage (emf) behind the source impedance to the value which would produce the specified –60 dBm level into a 150 Ω load resistor. Transducer gain is then computed as the dB difference between the measured terminated dBm output level of the device and the specified generator dBm output level.

Measuring Equalization and Emphasis

Many audio electronics devices have a deliberately non-flat frequency response. Equalizers are essentially sophisticated tone controls, designed to let the engineer create the desired tonal balance in the output signal. Both graphic equalizers and parametric equalizers are commonly used in the recording, sound reinforcement, and broadcasting industries. Frequency response of equalizers is normally measured in a conventional fashion by sweeping the frequency of an audio signal generator whose output level remains constant across the frequency band. The variations in output level resulting from various settings of the equalizer controls are normally the desired form for the information. Graphic and parametric equalizers typically have from 12 dB to as much as 20 dB of gain at some settings, so care must be taken to set the generator level sufficiently low so that clipping will not occur even at the points of maximum equalizer gain. Parametric equalizers may also have the ability to be set for a sharp rejection notch to eliminate one troublesome frequency. Notch depth may be 40 dB or greater. Measurement of a notch should be done with the generator level sufficiently high that noise in the equalizer under test does not become a limitation. Professional equalizers usually have a sufficiently-wide dynamic range that an input signal level around 0 dBu will neither cause clipping at maximum gain points nor noise limitation at minimum gain points.

Families of Equalization Curves

A desirable form of frequency response data display on equalizers is to plot a family of curves at various settings of the equalizer on the same graph. Figure 64 shows such a family for a graphic equalizer. Figure 65 is a family of several settings of the bass and treble tone controls of a consumer amplifier.


*Figure 64. Family of response curves, graphic equalizer.*

Figure 65. Family of response curves, bass and treble tone controls of consumer high-fidelity amplifier.

Loudness Controls

A special form of equalizer frequently found in consumer control preamplifiers, integrated amplifiers, and receivers is the loudness control. The loudness control concept is based on the fact that the frequency response of the human ear and brain is not constant at different sound pressure levels. As shown by the pioneering work of Fletcher and Munson, at very high sound levels, the human frequency response approaches flatness more closely. As the sound pressure level is reduced, the ear loses sensitivity more rapidly at low and high (bass and treble) frequencies than it does at midband frequencies. In an effort to partially compensate for this effect and produce perceptually-flatter response across a range of sound pressure levels, some consumer devices incorporate loudness controls. The loudness control adjusts equalization as it adjusts audio level. Some very basic loudness controls may simply switch in a fixed amount of bass and/or treble boost when the audio level is set below a certain point. More sophisticated level controls have continuous equalization functions dependent upon level, and in effect boost both low and high frequencies as they attenuate the overall audio level. Figure 66 is a family of response curves at several settings of a basic loudness control.


Figure 66. Family of curves, consumer amplifier loudness control.

Deviation From Standard Response

Devices such as RIAA phonograph record preamplifiers or tape recorder equalization circuits create a fixed equalization curve. Testing of such a device is normally done to determine whether it conforms to the standard specifications for that class of device within an acceptable tolerance. For this application, the best form of data presentation is a graph showing deviation from the standard curve. Such a "deviation from standard" curve can be obtained in one

of two ways with modern, automatic audio test equipment. One method involves a conventional response measurement followed by a computation. The computation subtracts the specified standard curve for the device from the actual measurement data, with the result being a graph of deviation from standard. Care must be exercised with this technique to keep the generator output level below the point which would cause overload at the frequency of maximum gain, which is normally not at a midband frequency such as 1 kHz. Thus, the common habit of setting the generator level for rated output at 1 kHz can lead to clipping at frequencies of higher gain.

Figure 67. Frequency response, RIAA-equalized phono preamplifier.

A second method involves dynamic control of test generator amplitude as a function of frequency during the sweep, with the amplitude controlled according to the inverse of the expected curve. For example, an RIAA-equalized phonograph preamp should follow a curve with a large amount of low-frequency gain and a large amount of high-frequency attenuation; see Figure 67. If the test generator follows the inverse function with low output at low frequencies and high output at high frequencies, a perfectly-equalized preamplifier would exhibit a constant output level at all frequencies. Any deviation from a flat output curve is then a measurement of deviation from perfect equalization. This second technique also helps avoid overload of the device during test, since the output amplitude will be approximately constant at all frequencies and a reference level set at midband should hold at other frequencies.

Emphasis

Preemphasis and deemphasis are terms applied to reciprocal values of fixed equalization used at two different points of a device or system, where the intent is that the two characteristics exactly offset one another to produce flat
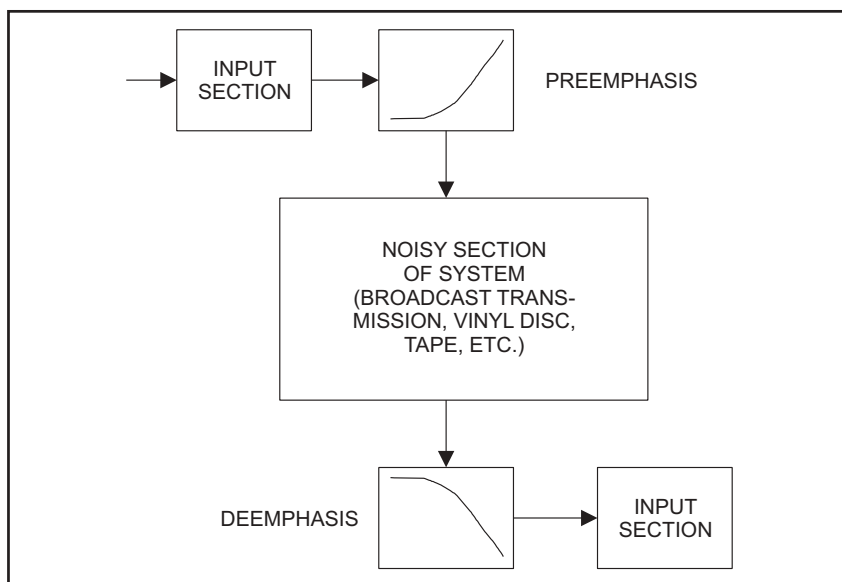
Figure 68. Block diagram, system with preemphasis and deemphasis.

overall response. The purpose of preemphasis and deemphasis is to improve the effective signal-to-noise ratio of the overall system. Noise introduced into a system at some point is often approximately white in its frequency distribution, resulting to the human listener in more noise at high frequencies. When it is known that a signal will be passed through a noisy section of a system, it is possible to boost the high frequency portion of the signal according to a specific curve before the noise-addition stage and then to pass the signal through a complementary filter which attenuates high frequencies after the section which introduces noise but before the listener hears it. If the two filter responses exactly match, the overall frequency response is perfectly flat but any high-frequency noise introduced between the filters is reduced by the attenuation of the second filter. The first filter is called the preemphasis filter and the second is called the deemphasis filter; see Figure 68.

Proper measurement of an emphasis system depends upon which portions of the system are being measured. If only the preemphasis circuit is being measured, this is a case of desiring to know deviation from standard response as described just above. The most suitable method is to cause the generator to vary according to the complementary (deemphasis) curve so that a perfect unit would display constant output. Deviations from this flat output are then easy to see. The same principle holds when measuring only the deemphasis portion; if the generator output follows the preemphasis curve as frequency is varied, the desired output response is flat.

If the entire system is measured end-to-end, the resulting overall response should be flat. This may thus appear to be a simple measurement and in fact is simple if there are no linearity limitations in the portions of the system be-

tween preemphasis filter and deemphasis filter. Commonly, however, those circuit blocks have an upper amplitude limit for linear operation. Since the preemphasis filter causes more output level at high frequencies, it is common, if care is not taken, to drive these portions of the system into saturation or clipping. Examples include FM broadcast systems, where the higher deviation resulting from preemphasis at high audio frequencies can exceed the deviation acceptance of the receiver. Analog tape recorders represent another similar situation, where high levels at high frequencies can saturate the tape. An end-to-end response measurement of a system including preemphasis and deemphasis should be made at a sufficiently low level that the linear operating range of the portion of the system between the emphasis blocks will not be exceeded. Good testing practice calls for response sweeps at several amplitudes. If the response varies at the higher amplitudes from the lower-amplitude curves, some maximum linear level is being exceeded. Typical preemphasis curves such as the 50 µs and 75 µs curves used in broadcasting produce 15 dB to 20 dB boost at the high end of the audio spectrum relative to 1 kHz (see the "Radios and Tuners" section below for more information and curves). Thus, a good starting point for generator level is about 20 dB below the specified maximum operating level of the system.

Spectrum Analysis

As noted in the "tools and techniques" section, hum or other interfering signals can create a false "floor" preventing measurement of true levels of distortion and noise. A spectrum analysis of the output of the device under test can be a good test of the integrity of the test setup—shielding, grounding, magnetic coupling from nearby transformers or CRT monitors, etc. Many of these effects are most likely to produce interference at the primary ac mains (power line) frequency of 50 or 60 Hz, depending on the power system in use at the test location. CRT monitors also often have large magnetic fields at their horizontal sweep frequency which may vary from about 15 kHz upwards, depending on the monitor design. A high-resolution spectrum analyzer such as an FFT analyzer can give the most accurate indication of the frequency of interfering signals, but even a swept one-third octave bandpass filter will give a good indication of such signals. When an FFT analyzer is used, a logarithmic horizontal display is desirable even though the FFT is an inherently linear process. A frequency-linear display lumps together all the potential low-frequency interfering signals at the left edge of the graph. A log display lets the operator separately measure the ac mains fundamental, 2nd harmonic, and 3rd harmonic. This is very useful since 2nd harmonic energy is an indication of inadequate filtering in a full-wave rectified dc power supply of the device under test, while strong fundamental and third harmonic are usually indications of magnetic coupling into the device or inadequate shielding.

FASTTEST Multitone Techniques

FASTTEST multitone techniques are directly applicable to all the audio devices in this category (amplifiers, equalizers, mixing consoles, etc.), and will produce greatly increased testing speed. Frequency response and interchannel phase response results will be identical between FASTTEST and stepped or swept single sine wave tests. Distortion and noise tests may produce somewhat different results due to differences in the technique.

Most audio devices in this category except for power amplifiers have a dynamic range extending from an upper clipping level (usually determined largely by the dc power supply voltages) to a noise level set by broadband noise. Distortion at levels below clipping is usually negligible compared to the wideband noise level, producing noise-limited distortion measurements at all usual operating levels. The absolute level of noise measured by FASTTEST in such devices should be identical to noise measurements by any other technique when differing measurement bandwidths are taken into effect.

The correlation between distortion measurements via FASTTEST versus conventional THD+N or IMD measurements is more difficult. First, the FASTTEST multi-tone signals are much more complex than THD+N or IMD test signals, and therefore produce more possibilities for intermodulation. Second, FASTTEST DISTORTION mode graphically plots distortion measurements into the portion of the audio spectrum where the distortion products actually occur, while a graph of THD+N versus frequency plots the amplitude of products of unknown frequency at the point in the spectrum corresponding to the sinewave stimulus signal which caused the distortion. Third, the relatively high crest factor of typical FASTTEST multi-tone signals means that less power is involved for any given peak signal level. That power is then further divided among many sinewave fundamentals instead of being totally concentrated at one frequency as with the single sinewave stimulus of THD+N testing. Fourth, THD+N is customarily stated either as a percentage of the sinewave signal level or in dB below that signal level. With the complex FASTTEST signal, it is not obvious what level should be used as a reference.

FASTTEST distortion and noise results are probably most useful when stated either in absolute units (dBu, etc.) or in relative dB with the DUT full-scale value being the 0 dB reference. This full-scale value must be determined from another measurement (such as clipping with a single sine wave), since it is not readily determined with the complex FASTTEST multi-tone signal.

If there are no distortion products in a portion of the spectrum, the FASTTEST DISTORTION mode measurements will exceed the FASTTEST NOISE mode measurements in that portion by exactly 3 dB since NOISE mode measures only half as many bins as are measured by DISTORTION mode. If the difference is more than 3 dB, this is an indication of intermodulation or harmonic distortion products falling into that portion of the spectrum.

Typical Performance, Consumer Preamplifiers	
Frequency Response 20 Hz–20 kHz	±0.1 dB to ±0.3 dB
Tone Control Range	±10 dB at 100 Hz and 10 kHz
THD+N @ Rated Output (80 kHz BW) 20 Hz–20 kHz	<0.01% to <0.1 %
S/N Ratio @ Rated Output from High Level Inputs (TUNER, CD, AUX)	80 to 100 dB
S/N Ratio @ Rated Output from Phono Inputs	70 to 80 dB
Stereo Separation 20 Hz–20 kHz	60 to 90 dB

Typical Performance, Power Amplifiers	
Frequency Response 20 Hz–20 kHz	±0.2 dB to ±1.0 dB
THD+N @ Rated Power (80 kHz BW) 20 Hz–20 kHz	<0.01% to <1.0 %
S/N Ratio @ Rated Power	90 to 120 dB
Stereo Separation 20 Hz–20 kHz	60 to 90 dB

Typical Performance, Graphic Equalizers	
Frequency Response 20 Hz–20 kHz at Flat Setting	±0.2 dB to ±0.8 dB
Maximum Boost or Cut of Any Band	±12 to ±15 dB
THD+N @ Rated Output (80 kHz BW) 20 Hz–20 kHz	<0.005% to <0.1 %
S/N Ratio @ Rated Power	80 to 100 dB

Typical Performance, Distribution and Line Amplifiers	
Frequency Response 20 Hz–20 kHz	±0.05 dB to ±0.1 dB
THD+N @ Rated Output (80 kHz BW) 20 Hz–20 kHz	<0.003% to <0.01 %
S/N Ratio @ Rated Output	100 to 120 dB

Companders

Compressors, expanders, limiters, clippers, modulation level and loudness maximizers, and noise reduction units are all grouped into the general category of companders. These are real-time electronic audio devices such as discussed previously, but they introduce a new factor. All these devices have deliberately non-linear input versus output curves; see Figure 69. A portion of their characteristic curves is usually linear, where 1 dB input change produces 1 dB output change. Above or below some threshold (often called the "knee" of its characteristic), they show their compression or expansion characteristic where 1 dB of input change results in less or more than 1 dB output change. The most sophisticated of such units may first divide the audio frequency spectrum into a number of frequency bands, then apply compression or limiting with different compression ratios or time constants in each band, and finally recombine the bands into a single output. Companders require some additional testing considerations due to their non-linearity, and particularly in measuring the dynamics of that non-linearity.

Companders may be tested both statically and dynamically. The standard battery of tests is static; measurements are made only after the DUT has stabilized under the new stimulus conditions at each step of a sweep. All the conventional audio measurement techniques may be used on a compander. The inherent characteristics of the device will cause different results with some of these measurements. For example, frequency response measurements made at amplitudes well below the compression knee will behave the same as on a linear amplifier. If the signal is at or above the compression threshold and if there are response variations in the DUT before the circuit point where the compressor senses level, the compression action will suppress or eliminate these variations in frequency response.

Noise Measurements on Companders

Conventional techniques of noise measurement may give surprising answers when used to measure companders. For example, when making the signal portion of a signal-to-noise ratio measurement on a compressor with a reference signal well above the threshold of compression, the compressor will reduce its gain significantly. When the signal is removed to measure noise, the compressor will increase gain and the resulting S/N ratio will be lower than would be expected on a linear amplifier. The resulting noise is not typical of the actual noise added during normal operation. If the compressor has a relatively long release time, it will take significant time for the noise measurement to stabilize when signal is removed.

As noted earlier in the FASTTEST discussion of the "tools and techniques" section, the FASTTEST method can measure noise in the presence of test signal. Noise is measured in "empty bins" all across the spectrum. These empty

bins alternate with bins possibly containing generator fundamentals or distortion products. If the amplitudes of the individual sine waves making up the FASTTEST signal are selected to produce the same spectral energy distribution as the types of program material normally handled by the compressor, the gain of the various frequency bands of the compressor will be set to the typical operating levels which they would have with actual program material. The FASTTEST NOISE function then measures noise under normal operating conditions.

Unique Tests for Companders

The unique characteristics of companders make it desirable to add some tests to the standard battery of tests. Figure 69 is a static input/output linearity test. The generator output level (at a fixed frequency) is swept across a wide dynamic range and the corresponding device output level is measured and plotted on the vertical scale by the analyzer. At each level, the DUT was allowed to stabilize before the output level measurement was made. The break point in the curve shows the onset of compression or limiting. The compression ratio above the knee is given by the slope of the graph, which is approximately 0.5 dB output change per dB input change in this example. A compressor may have more than one knee and several different compression ratios.

AUDIO PRECISION   STEADY-STATE INPUT-OUTPUT LINEARITY, COMPRESSOR

Figure 69. Steady-state output vs input, compressor.

The dynamic characteristics of a compressor may be measured in both the time domain and frequency domain. The typical signal for dynamic tests is a sine wave burst. Multi-band compressors can be tested at various points across the spectrum by changing to different frequency sine waves in the burst. Typi-

AUDIO PRECISION   COMPRESSOR OUTPUT, 50 ms BURST WITH 20 dB AMPLITUDE RATIO

Figure 70. Dynamic testing of compressor with burst signal.

cally, the sine wave is not reduced to zero amplitude between bursts, but instead operates at a calibrated lower amplitude such as –20 dB. This two-level signal is a simple simulation of peak and average levels of voice or music.

To measure dynamic performance in the time domain, it is desirable to display the envelope of the burst at the compressor output rather than the sine waves within the burst. If the measuring instrument has the capability of displaying the burst amplitude logarithmically or in decibels rather than linearly, it is much easier to study the device performance. The key elements to be determined in time domain analysis are the duration and amplitude of any overshoot at the leading edge of the test burst (before the compressor has time to exercise control and reduce the output to the target value), the amount of gain reduction during the burst, and the duration and shape of the gain recovery following the burst. The time interval between bursts should be long enough that full gain recovery occurs before the following burst. Figure 70 is an example of a burst envelope at the output of a compressor. Overshoot occurs during the first 5 ms of the burst. Output level is then controlled to 10 dB above the pre-burst level, even though the height of the stimulus burst is 20 dB. At the end of the burst, the output immediately falls by 20 dB since gain is still in the reduced state, then gradually returns to normal during the next 150 ms.

Frequency domain dynamic analysis is done by FFT spectrum analysis. By performing the FFT on a short record length, it is possible to measure distortion at various points during the attack, steady-state gain reduction, and decay portions. For example, at the 48 kHz sample rate typical of audio applications,

AUDIO PRECISION   FFT, OVERSHOOT PORTION OF BURST (0-5 ms)

AUDIO PRECISION   FFT, MAXIMUM COMPRESSION PORTION OF BURST (12-17 ms)

AUDIO PRECISION   FFT AFTER FULL GAIN RECOVERY (190-195 ms)

Figure 71. FFT Compressor output during attack/overshoot, steady-state compression, and gain recovery phases of burst signal in Figure 70.

a 256-point sample has a time duration of 5.33 ms. The short record lengths required for good time selectivity produce lower frequency resolution. The frequency resolution (bin width) for this example is 187.5 Hz. If the sine wave burst is at a frequency of 1 kHz, for example, this resolution is still adequate to measure distortion down to levels of 60 dB to 80 dB since harmonics are spaced more than 10 bins apart. Figure 71 a, b, and c are examples of distortion during the attack (overshoot), steady gain reduction, and gain recovery portions of the compressor characteristic shown in time domain form in Figure 70. It can be seen that 3rd harmonic distortion remains below –55 dB to –60 dB during the steady gain reduction and gain recovery portions, but is about –48 dB during the initial overshoot at attack.

Tape Recorder Alignment and Performance Measurements

Although digital recording techniques have many advantages and are gaining market share, analog tape recorders still dominate professional and consumer audio applications all over the world. Analog tape recorders have some unusual test requirements not discussed in the previous sections. First, obtaining optimum performance from analog tape recorders requires alignment which must be done in conjunction with quantitative measurements for best results. Secondly, many common performance verification tests must be performed somewhat differently on tape recorders due to the wow and flutter which the recorders introduce, and due to the time delay caused by tape transit between record and playback heads in simultaneous record-reproduce mode.

Alignment

Nearly all professional audio tape recorders and the better consumer recorders use separate heads for recording and reproduction (playback). Both mechanical and electrical alignment are required for best performance from a tape recorder. Furthermore, there is an interaction between mechanical and electrical alignment of the recording head. The adjustments commonly required during alignment are:

- Reproduce head azimuth adjustment.

- Reproduce level adjustment (midband).

- Reproduce high-frequency (treble) level adjustment (equalization).

- Recording head azimuth adjustment.

- Recording level adjustment (midband).

- Recording bias level adjustment.

- Recording high-frequency (treble) level adjustment.

Some tape machines may also have low-frequency (bass) level adjustments in recording and/or reproduce sections.

Reproduce Section Alignment

All reproduce section alignment and adjustments are made while playing a reference tape. Normally reference tapes are purchased from a company which individually records them on extremely-well aligned high quality tape machines.

Head azimuth adjustments are made in order to orient the magnetic head gap perpendicular to the direction of tape travel. Stereo tape machines com-

prise the large majority of tape machines in the professional audio and broadcast fields. On these machines, reproduce head azimuth is most accurately aligned while comparing the phase relationship between the two channels when playing a high-frequency signal on a reference tape recorded on a perfectly-aligned machine, often by a monaural (full-track) head. If azimuth is not correct, the signal arrives earlier at one channel of the head than the other and a phase difference is measured. When azimuth is properly aligned, there will be a zero average phase difference between the two channels of the head. Multi-track tape machines can be similarly aligned by using the two outside edge tracks for a phase measurement. Signal frequencies of 8 kHz to 16 kHz are typically used for azimuth alignment, since the sensitivity of the variation in phase versus azimuth angle error increases with increasing frequency. However, this same sensitivity also means that it is easier at high frequencies to accidentally mis-align for an apparent zero phase reading which is really 360 degrees or 720 degrees, with consequent azimuth error. For this reason, common practice is to first do a rough azimuth alignment using a relatively low frequency on the tape and then move to a higher frequency for a final, more precise adjustment. If a reference tape is available with the multi-tone FASTTEST signal recorded, the results are graphed as phase versus frequency and a flat zero degree curve at all frequencies is the desired result. The 360/720 degree error at high frequencies cannot occur since such an error would be accompanied by large errors at all the other high frequencies and the curve would not be flat at zero degrees.

Reproduce level adjustments at mid-band and high frequency are made while playing appropriate sections of a reference tape. The mid-band reproduce level adjustment is made while playing the reference fluxivity section of the reference tape, and the level controls on each channel are adjusted for the specified reference output level from the tape machine. Not all reference tapes are recorded at the same reference fluxivity level, so a correction will be necessary if the tape machine is being aligned for a different reference fluxivity than that recorded on the reference tape. For example, if the tape machine is to be set up for a +4 dBu output level with 250 nW/m fluxivity but the available reference tape is recorded with 320 nW/m fluxivity, the reproduce level should be set up to produce an output level 320/250 or 2.2 dB higher than the target (+6.2 dBu rather than +4.0 dBu). The frequency response portions of the reference tape are often recorded at levels lower than the reference fluxivity section to avoid saturation problems. This is particularly likely to be true with reference tapes, at 7.5 inches per second (ips) (19 centimeters per second) and slower. These frequency response signals are typically either 10 dB or 20 dB below reference fluxivity; this amount will be indicated in the documentation furnished with the reference tape and may also be stated by voice narration on the reference tape. Note that reference tapes which have been used many times will lose signal amplitude at high frequencies before their characteristics change at mid-frequencies. If there is any question on the quality of the reference tape, buy a new one.

If the tape machine includes a low-frequency (bass) level adjustment in the reproduce section, this adjustment cannot be properly made while playing the usual full-track (monaural) reference tape. At low frequencies, a magnetic fringing effect occurs at the edges of the tracks. Unless a reference tape recorded on a stereo tape machine is available, the reproduce bass level control cannot be adjusted until later, during alignment of the record section.

Record Section Alignment

On three-head tape machines which can simultaneously record and reproduce, all record section alignment is done in record-reproduce mode, using a blank tape. The complete reproduce section alignment (with the possible exception of low-frequency level, as described in the previous paragraph) must be completed before the record section alignment is done, since any errors in the reproduce section will directly affect record section alignment. It might seem that reproduce section errors can be compensated for by record section adjustments, but the result is unacceptable when tapes must be transferred between machines.

The test system used during record section alignment must be capable of inserting a time delay after each frequency change at the recording head before measuring phase from the reproduce head, to allow for the tape transit time from record to reproduce head. The minimum time delay may be computed from the tape speed and the spacing between record and reproduce heads.

Record level adjustment is done by sending a calibrated reference level from the audio generator to the machine input (often +4 dBu) and adjusting the record level controls of both channels for reference output. If the level indicators (VU meters) built into the machine are properly calibrated, this will also correspond to a 0 VU indication.

The record head azimuth on a three-head machine must be aligned to match the reproduce head. Measurements are made of inter-channel phase from the reproduce head while in simultaneous record-reproduce mode. The adjustment could be made while recording a single high frequency tone on both tracks. The possible 360 degree error described earlier can be avoided by more sophisticated techniques such as repeatedly stepping the test tone through several frequencies or using the multi-tone signal of the FASTTEST technique. For sweeping, a three-frequency repeating sweep with the frequencies spaced logarithmically from 1 kHz to 16 kHz is a good choice. If phase can be graphed versus frequency on a continuously redrawn graph, the goal is to obtain a flat horizontal zero degree line at all three frequencies. Misalignment of azimuth will result in the curve climbing or drooping at high frequencies, as shown in Figure 72. The FASTTEST technique produces a similar graphic display. The final flat curve proves that proper alignment has been made rather than misalignment with a 360 degree error at some high frequency, since such an error would also produce non-zero readings at the lower frequencies. Accu-

AUDIO PRECISION   PHASE vs FREQUENCY AT SEVERAL AZIMUTH SETTINGS

Figure 72. Graphs of phase vs frequency for several azimuth settings of tape recorder head.

racy within several degrees is acceptable. Sweep-to-sweep phase jitter is likely to vary from a few degrees to ten or fifteen degrees at the highest frequencies, depending upon machine quality.

Recording bias level adjustment is perhaps the most important and most difficult in the machine. Bias is a high-frequency ac signal, typically at a frequency three to six times the maximum response frequency of the tape machine. Bias is combined with the audio signal to be recorded in order to optimize the frequency response, distortion, and noise characteristics of recording. Different types of magnetic tape formulation require different amounts of bias for optimum operation, and any given tape type will also require different amounts of bias at different tape speeds. Different batches of the same nominal tape formulation may require slightly different values of bias for best results. Thus, a broadcasting or recording studio is likely to be re-aligning tape machine bias many times.

There are several schools of thought on the best method of optimizing bias. The most common techniques are to adjust bias for minimum distortion, for minimum modulation noise, or by the overbias technique using manufacturer-furnished numbers. The overbias technique is the most common.

The minimum distortion technique requires recording at a high level (perhaps +3 to +10 VU) with a mid-frequency signal and adjusting bias level for a distortion minimum at the machine output. Near the distortion minimum, there may be a rather rapid change in output level versus bias level. The minimum

distortion technique is best practiced with an audio analyzer which can simultaneously measure both distortion and output level from the tape. Bias is first adjusted to the approximate point of minimum distortion. Then, watching both distortion and output level, a final adjustment of bias is made to produce the highest output level without a significant increase in distortion.

The minimum modulation noise technique is nearly an art and requires more experience than the other techniques. It is commonly practiced while recording a very low-frequency tone, such as 30 to 50 Hz, at a high level. The operator listens to the output of the tape machine while adjusting bias for minimum modulation noise. The recorded tone is so low in frequency that it not reproduced through the small monitoring loudspeakers typically built into tape machines. The art comes in separating modulation noise from other types of noise, and no objective measurement techniques have been developed to do this reliably. Experienced operators are able to detect a character to this noise and adjust bias to minimize it.

The overbias technique is a method of systematically arriving at an optimum compromise bias level as pre-determined by either the magnetic tape manufacturer or tape machine manufacturer. Tape and tape recorder manufacturers will typically furnish an overbias value for all common types of professional tape and at all commonly-used tape speeds. For example, a manufacturer may specify 1.5 dB overbias at 30 and 15 ips, 4.5 dB at 7.5 ips, and 6 dB at 3.75 ips. The overbias value refers to the amount which the bias should be increased above the value which produces maximum reproduced level recorded with a low-amplitude high-frequency signal. 10 kHz is the recording frequency commonly used for overbias adjustments. An audio signal typically 20 dB below normal (–20 VU) is sent to the tape machine from the audio generator and bias is set to minimum. As in all record-reproduce adjustments, time must be allowed for the tape to move from record to reproduce head before the effect of the changed bias level can be seen. The bias level is slowly increased while the reproduced output level is observed. Bias is continually increased while the output level increases until the bias value is found which causes peak output. The output level will start to drop off at still-higher values of bias. The output level at this peak is noted. Bias is then further increased until the output level drops below the peak value by the specified overbias value for the tape type and speed in use.

There is a small amount of interaction between bias level and sensitivity (output level) even at midband frequencies such as 1 kHz. Thus, it is desirable to check and possibly re-adjust the record level after setting bias. There is greater interaction between effective azimuth angle of the record head and bias. The effective azimuth angle is a combination of the mechanical azimuth angle of the tape head and the amount of bias current flowing through the head. Thus, an iterative process may be required of adjusting record head azimuth, adjusting bias, re-adjusting azimuth if necessary, re-setting bias, etc.

This iteration may not be necessary after a machine is first properly aligned, with only small "touch-up" changes made in bias thereafter.

Record treble level is set after bias has been aligned. Tape recorder alignment instructions commonly call for sending a specific high frequency signal (whose frequency depends upon the tape speed being aligned) and adjusting for the same output level obtained at the mid-band reference frequency. Treble level adjustments are normally made 20 dB below reference fluxivity level to avoid tape saturation problems, although at 15 ips and faster saturation should not be a problem. A better overall compromise between response flatness and high-frequency response extension may be made by adjusting treble level while making a repeating frequency sweep. If the audio test set permits a repeated step through an arbitrary set of frequencies, a frequency set like this can be useful:

<div align="center">

1 kHz
10 kHz
12 kHz
15 kHz
18 kHz
20 kHz

</div>

The high-frequency response curve may then be continuously re-graphed while adjusting treble level for the best compromise between flat response and extension to the highest possible frequency.

Finally, low-frequency (bass) level adjustments can be made while using a repeating sweep at low and mid frequencies. A set of frequencies such as:

<div align="center">

30 Hz
50 Hz
70 Hz
100 Hz
200 Hz

</div>

may be a good starting point.

Two-Head Tape Recorders

Not all tape recorders have separate record and playback heads, with the ability to reproduce from the tape as it is recorded. Some two-head tape recorders are playback-only machines. Common examples include consumer personal cassette players, automobile cassette players, and broadcast tape cartridge machines. Other tape machines, such as most home-type consumer cassette machines, can both record and reproduce tapes but not simultaneously, since the same head is used for recording and for reproduction.

The reproduce sections of all two-head machines are aligned exactly as are the reproduce sections of three-head tape recorders described earlier. A reference tape is required, and head azimuth, reproduce level, and reproduce treble level are aligned as described earlier. Complete reproduce frequency response measurements can be made if the reference tape contains the necessary signals across the audio band.

For two-head machines which can record, the record section alignment process is generally a good deal more cumbersome and awkward than on three-head machines. The only exception is head azimuth; since the same head is used for playback and record, the azimuth adjustment is not touched following alignment with a reference tape. All other record section alignment steps necessarily involve recording a section of tape at preliminary settings for the parameter being aligned, rewinding the tape, playing it back and measuring during playback to see if the preliminary setting is correct or if further adjustment is needed. The most efficient procedure usually involves recording consecutive sections of a blank tape at several settings of the parameter being adjusted, while keeping a log of the setting versus the tape location counter. Then, when the tape is reproduced, the resulting output measurements versus tape counter reading are noted and the most appropriate setting restored following all the measurements. For example, the bias control could be adjusted to ten progressively higher settings for five seconds each as the cassette is recorded with a 10 kHz –20 dB signal. During playback, the output level is logged as a function of time. If the ten settings included the appropriate range, the resulting amplitude function of time should tell which bias control setting is closest to the desired overbias value. If necessary, a second recording may be required with smaller parameter adjustments near the optimum value to arrive at the most accurate setting. Record level and record treble level can be similarly set.

Performance Measurements

After the tape recorder has been aligned, performance measurements should be made to determine if the overall performance is acceptable. The reproduce section frequency response can be measured and graphed if the audio analyzer has the capability of measuring and graphing from an external signal consisting of a progression of tones. Since many reference tapes have voice announcements before each tone, the analyzer must be able to ignore the voice and measure only the tone in order to produce a clean graph. Excellent professional analog tape machines can produce reproduce frequency response within ±1 dB from 300 Hz to 18 kHz or 20 kHz at 7.5 ips, extending (but at perhaps ±2 dB) to 22 to 25 kHz at 15 ips and 25 to 30 kHz at 30 ips. Good-quality consumer cassette machines can produce flat response within ±2 or ±3 dB from low frequencies to 20 kHz. As noted earlier, the low frequency response of a stereo tape machine cannot be properly measured using the normal full-track reference tape. Total harmonic distortion or third harmonic distortion of the midband reference fluxivity signal will normally be well below 1%.

Indicated values of THD+N measured from a tape machine (or phonograph turntable) may be influenced by wow and flutter, depending upon the specific architecture of the distortion analyzer. Flutter produces frequency modulation of all recorded tones, resulting in a family of FM sidebands around the tone. THD+N analyzers which use a narrow notch filter to reject the fundamental tone may not totally reject the FM sidebands associated with flutter of that tone. If the sidebands are not totally rejected, they influence the THD+N value along with true harmonics and wideband noise. Some analyzers have more so-phisticated notch filters (usually four-pole) whose notch shape is wider with a more square bottom. Properly designed, such notches will also reject flutter sidebands. It is also possible to follow the notch filter with a high-pass filter with zero attenuation at second harmonic and above but considerable rejection near the fundamental tone and below to attenuate flutter sidebands. If individual harmonics are measured, the selectivity of the measurement technique rejects flutter sideband energy.

Figure 73. Tape recorder frequency response at three speeds.

Combined record-reproduce frequency response can be flat within ±1 dB from 30 Hz to 18 kHz at 7.5 ips. At the higher speeds, high-frequency response will be extended as noted above for reproduce section response, but the low-frequency cutoff point also increases. Thus, a given tape machine may be flat down to 10 Hz at 7.5 ips, 20 Hz at 15 ips, and 35 Hz at 30 ips. Figure 73 shows overlaid graphs of record-reproduce frequency response on a moder-ately-priced professional tape recorder at three speeds.

Signal-to-noise ratio improves with faster tape speeds and wider tapes. Thus, 20 kHz bandwidth noise levels (without noise reduction) may be around

50 dB with cassette tapes, improving to the upper 50s for 1/4 inch tape at 7.5 ips and approaching 70 dB at 30 ips. Third harmonic or total harmonic distortion at 0 VU should stay below 1% from 100–200 Hz up through the entire audio spectrum, but will rise at low frequencies to perhaps several percent. Most well-aligned professional tape machines will be able to record midband frequencies at levels of +9 to +12 VU before exceeding 3% distortion.

Crosstalk varies strongly with frequency in stereo tape recorders. At low frequencies (below 100 Hz), crosstalk may be as bad as 15 to 20 dB due to coupling of the long wavelength signals between the magnetic heads of the two channels. At midband frequencies, crosstalk is commonly better than 50 dB. At very high frequencies, crosstalk may again degrade to 30 dB or less.

Wow and flutter values are generally lower at the faster tape speeds. Wow and flutter is also typically lower on reel-to-reel machines than on cassette decks. Values below 0.1% (IEC/DIN, weighted) are reasonable on good quality consumer cassette decks, while high-quality professional tape machines sometimes approach 0.01%.

Typical Performance, Analog Cassette Recorders	
Frequency Response 50 Hz–18 kHz	±2 dB to ±3 dB
THD+N @ +3 VU, 400 Hz, 22 kHz BW	1% to 3%
S/N Ratio @ +3 VU with noise reduction OFF	50 to 55 dB
S/N Ratio @ +3VU with noise reduction ON	55 to 65 dB
Stereo Separation 400 Hz–10 kHz	35 to 50 dB
Wow & Flutter (IEC/DIN Weighted)	0.05% to 0.2%

Typical Performance, Professional Reel-Reel 1/4" Tape Recorders	
Frequency Response 35 Hz–20 kHz	±1 dB to ±2 dB
THD+N @ 1 kHz +3 VU (22 kHz BW)	1% to 3%
S/N Ratio @ Rated Output	60 to 70 dB
Stereo Separation 400 Hz–10 kHz	45 to 60 dB
Wow & Flutter (IEC/DIN Weighted)	0.01% to 0.07%

FASTTEST

As noted in the alignment section above, the multi-tone FASTTEST signal can be used for all record section alignment other than bias setting and can

also be used for reproduce section alignment if suitable reference tapes are available.

When the FASTTEST technique is used for performance measurements, several factors must be considered and specific FASTTEST features must be used. Flutter in a tape machine causes FM sidebands around each sine wave signal. At higher frequencies, these sidebands will fall into FFT bins above and below the bin which the sine wave fundamental occupies. In an FM process, all the energy which goes into the sidebands comes out of the "carrier." For response measurements, FASTTEST normally measures only the amplitude of the bin which is expected to contain the generated sine wave fundamental. As the energy going into flutter sidebands becomes significant, the energy remaining in each sine wave fundamental bin is reduced and the effect is as if high frequency response is rolling off. Another selectable response mode of FASTTEST specifically incorporated for tape machines compensates for this effect. It combines the amplitudes of all FFT bins within a specified frequency resolution of the nominal frequency in an RSS (root sum square) fashion and reports the resulting amplitude. This is mathematically equivalent to re-combining the flutter sideband energy back into the "carrier" signal, and the result will be a proper frequency response measurement even in the presence of wow and flutter.

The FASTTEST technique requires that all sine waves in the signal to be analyzed be synchronous with the FFT analyzer buffer. When a reference tape is recorded with a FASTTEST multitone signal on one analog tape recorder and later reproduced on another, there will be sufficient difference in speed of the two machines to eliminate the synchronous relationship. The FASTTRIG variation on FASTTEST includes the additional feature of frequency error correction. FASTTRIG uses a digital copy of the expected waveform as a reference signal and corrects all frequencies in the measured signal to achieve the synchronous relationship. All alignment and performance measurement of reproduce sections of analog machines with a FASTTEST multitone reference tape must therefore be done using the FASTTRIG program and with a downloaded multitone reference signal identical to the signal originally recorded onto the reference tape.

DISTORTION mode in FASTTEST accumulates (by an RSS computation) the amplitude of all bins between each adjacent pair of sine waves in the test signal. This value thus represents the total energy of all harmonic and intermodulation distortion products and noise falling into that portion of the spectrum. When a tape recorder adds flutter sidebands to each tone, those flutter sidebands falling outside the central bin of each tone would be treated as if they were distortion and noise by DISTORTION mode. The result will be an apparent large increase of distortion and noise at higher frequencies. To eliminate this problem, the bins surrounding each tone up to the specified frequency resolution will be excluded from the DISTORTION summation. The result will be a more accurate indication of true distortion and will correlate better

with human perceptions, since psychoacoustic masking mechanisms prevent close-in sidebands from being audible.

NOISE mode of FASTTEST cannot be used with tape recorders having any significant amount of wow and flutter, or any significant speed error between recording and playing a tape. NOISE mode depends upon alternating "empty bins" typically about 2.9 Hz in width interleaved between bins which may contain generated signals or distortion products of those signals. Flutter and/or speed error will move enough energy into those theoretically-empty bins to prevent NOISE function from being useful.

Radios and Tuners

Two key factors make measurement of the audio sections of radio tuners and receivers different from the more common real-time audio devices discussed previously. First, the input to such devices must be a suitably-modulated radio-frequency carrier signal rather than a direct audio input. Second, many broadcasting systems include preemphasis in the transmission circuits and matching deemphasis in the reception circuits. The purpose of the receiver deemphasis is to reduce the effects of noise in the transmission path, since white noise has most of its power at high audio frequencies. Matching preemphasis is then required in the transmitter in order to produce overall flat response. The implications of these emphasis circuits must be taken into effect in measurements.

Audio program transmission at medium-wave frequencies (approximately 500 kHz to 1600 kHz) is accomplished by amplitude modulation (AM) techniques. Audio transmission at the very-high frequency (VHF) bands (88 MHz–108 MHz in most of the world, slightly lower in Japan) is accom-
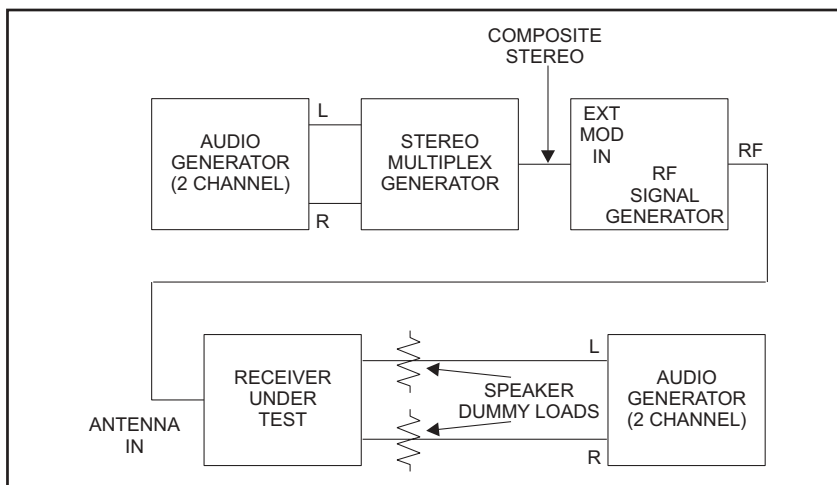
Figure 74. Test equipment block diagram, radio receiver test system.

plished by frequency modulation (FM). The audio channels of TV broadcasting also use FM. A necessary instrument for testing radios and tuners is thus an RF signal generator covering the necessary range with modulation capability of the necessary type. If the receiver is a stereo device, the RF signal generator must also be used with an appropriate stereo multiplex generator to convert left and right channel inputs from an audio generator to a composite stereo signal to feed to the RF signal generator modulation input. See Figure 74 for a simplified block diagram of the necessary instruments and connections.

Tuners typically have nominal audio output levels on the order of 1 V to 2 V rms, designed to work into load impedances of 10 kΩ or higher. Thus, the tuner output may be connected directly to the audio analyzer input. Stereo tuners have left and right channel outputs which connect directly to the A and B channel inputs on two-channel analyzers. Receivers may supply an intermediate output connector similar to tuner outputs, but definitely will have audio power amplifier output connections. These connections are designed to work into loudspeaker load impedances, typically at 8 Ω but often lower in automotive stereo applications. These outputs should normally be connected to power resistors of the correct resistance value and capable of safely dissipating the rated output level of the receiver. Any tone or loudness controls in the receiver should be disabled or put into their flat response position before making measurements on the receiver.

Prior to approximately 1989–1990, AM broadcast transmissions did not involve any preemphasis at the transmitter. Many individual broadcasters were deliberately using equalization to boost high frequencies prior to the transmitter to help compensate for the severe high-frequency loss of typical inexpensive AM receivers. This high-frequency boost was formalized in the US. in the NRSC preemphasis standard, which follows a 75 µs preemphasis curve up to approximately 10 kHz with no further boost above that frequency. There is a rather ambiguous situation as to whether AM receiver testing should be done with this preemphasized signal, since this preemphasis is not widely used outside the US.

There is no ambiguity in the use of emphasis in FM and TV aural broadcasting, however. North American and Japanese broadcasting standards use 75 µs preemphasis in FM and TV aural transmitters, and most of the remainder of the world uses 50 µs preemphasis. The 50 µs and 75 µs figures refer to the time constant of a capacitor-resistor single-pole high-pass filter circuit which will produce the desired preemphasis. Figure 75 shows the filter and Figure 76 shows the resulting response for 50 µs and 75 µs time constants. Receivers and tuners are then expected to use the corresponding deemphasis circuit in their early audio stages, following the FM detector. An audio-frequency sweep at a low amplitude (10% deviation or less at 1 kHz) into an actual preemphasized broadcast transmitter would result in flat frequency response at the deemphasized outputs of a receiver or tuner, since the transmitter and receiver characteristics should exactly offset one another.
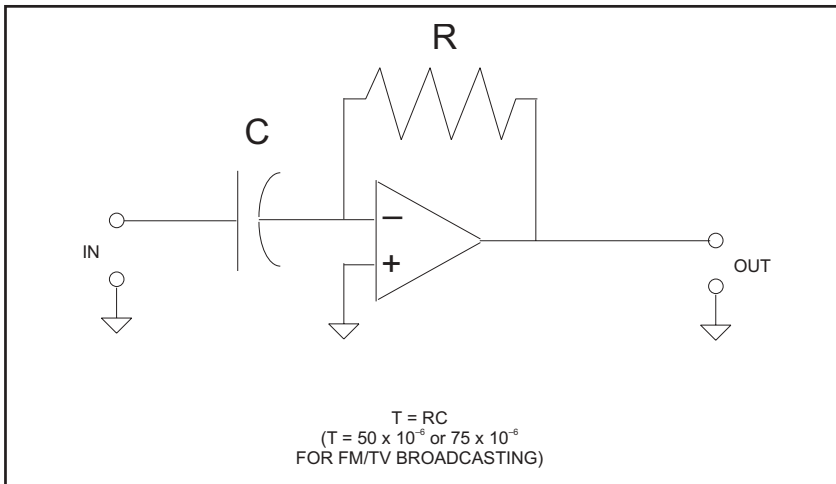
Figure 75. Equivalent filter circuitry, simple preemphasis network.

   RF signal generators often do not have preemphasis circuitry built in. With a constant-amplitude audio frequency sweep into their external modulation terminals, they produce constant deviation at all audio frequencies. This mode of operation is convenient in determining whether the receiver or tuner will accept full-rated deviation at all frequencies without excessive distortion. Measurements of THD+N versus frequency should normally be made at constant, maximum deviation.
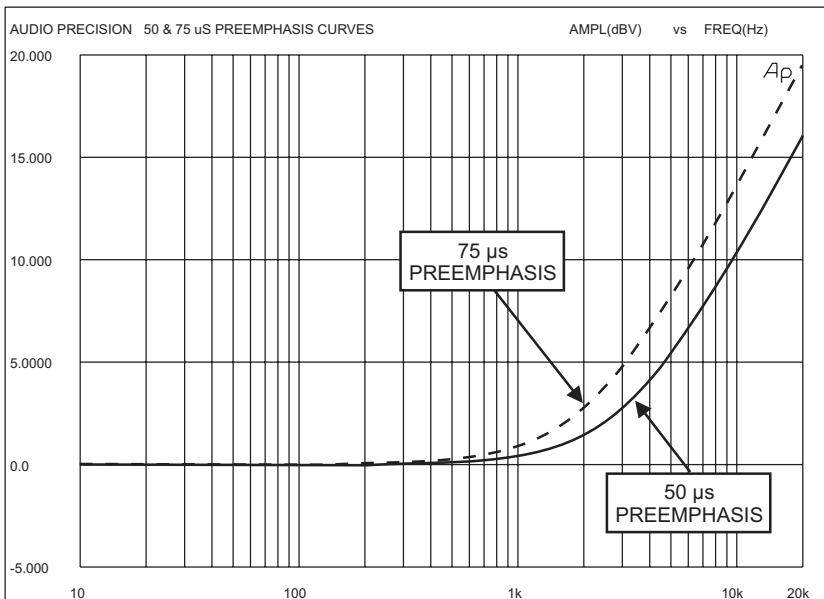
Figure 76. Frequency response curves, 50 μs and 75 μs preemphasis.

Operation at constant deviation will not, however, produce flat frequency response from the receiver. It will instead show a deemphasis curve response. Audio analyzers which have the ability to subtract a stored reference file from measurement results can correct this deemphasis curve to a nominally flat response, with any deviations from perfect flatness then being an indication of error in the receiver's deemphasis circuitry. Alternately, some audio test sets have the ability to automatically adjust the audio generator output amplitude according to some function of frequency as frequency is varied. If a standard preemphasis curve is used as the reference file for an audio sweep from such a generator connected to the external modulation input of an RF generator, the resulting sweep should produce flat response from the deemphasized receiver or tuner. The deviation level at the highest modulating frequency (typically 15 kHz) dare not exceed the maximum rated frequency deviation of the receiver. Variations from flatness are measurements of errors in the receiver deemphasis circuitry. High-quality tuners and receivers can deliver frequency response within ±1 dB from very low frequencies to approximately 15 kHz, but response is often starting to roll off very rapidly near 15 kHz in order to get sufficient attenuation of the stereo pilot tone at 19 kHz. Distortion of high-quality tuners can be below 0.1% (–60 dB) at full deviation. Stereo separation on the best tuners and receivers can reach 50 to 70 dB, although anything better than 30–40 dB is probably unnecessary for full stereo effect.

In principle, noise measurements of tuners and receivers are made in the same fashion as amplifiers. At a high RF signal amplitude which produces full quieting of the receiver, the audio generator output is removed and the resulting noise level from the receiver is measured. In practice, many RF signal generators have such high levels of residual frequency modulation (incidental FM) that the specifications of better-quality tuners and receivers (–70 dB and better) cannot be verified. High-stability quartz-crystal-controlled generators may have sufficiently low residual FM to make definitive noise measurements.

SINAD

A measurement unique to FM receivers is SINAD; the first syllable is pronounced exactly as the English word "sin." SINAD is an acronym for:

$$\frac{(Signal + Noise + Distortion)}{(Noise + Distortion)}$$

SINAD is measured with a conventional THD+N analyzer using the notch filter technique, and in fact SINAD is merely the reciprocal of THD+N, stated as a positive ratio. Thus, 3.3% THD+N (–30 dB THD+N) corresponds to a 30 dB SINAD. SINAD is typically measured with either a 400 Hz or 1 kHz modulating signal, adjusted in amplitude to produce 2/3 of maximum rated deviation of the system. With 75 kHz deviation corresponding to 100% modulation for FM broadcasts, SINAD would thus be measured at 50 kHz deviation.

AUDIO PRECISION   THD+N (dB) vs ANTENNA TERMINAL SIGNAL (dBf)

Figure 77. Typical graph of SINAD vs RF signal strength, FM receiver.

SINAD is used to state the sensitivity of FM receivers and also in measurements of adjacent channel and alternate channel selectivity and other interference rejection measurements. The SINAD technique was developed to replace the earlier "quieting" method of specifying receiver sensitivity. With the quieting method, all RF signal was removed (and the squelch circuit defeated if present) and a reference level measurement was made on the resulting audio noise output of the receiver. The RF signal generator amplitude was then brought up from zero until the noise level was reduced ("quieted") by the specified number of dB, often 20 dB. The problem with the quieting-based sensitivity specification was that it encouraged receiver designers to make the receiver bandwidth very narrow. This could produce a receiver with impressive quieting sensitivity, but which distorted on program material at high deviation levels. Since SINAD is measured with a modulation value typical of average program levels, narrowing the receiver bandwidth increases distortion even though it reduces noise and designers are thus not encouraged to produce unrealistically-narrow bandwidths. Figure 77 shows a typical curve of output SINAD versus RF signal level at the receiver antenna terminals.

Typical Performance, FM Stereo Tuners and Receivers	
Frequency Response 50 Hz–13 kHz	±0.5 dB to ±2 dB
THD+N @ 100% deviation 1 kHz (22 kHz BW)	.05% to 1%
Monaural S/N Ratio @ 100% deviation with strong RF signal	60 to 80 dB
Stereo S/N Ratio @ 100% deviation with strong RF signal	50 to 60 dB
Stereo separation @ 1 kHz 100% deviation with strong RF signal	30 to 50 dB

FASTTEST in Receiver Testing

The multi-tone FASTTEST technique can be used for rapid testing of receivers and tuners, as is typically required at the very high production volumes typical of auto stereo manufacturing. A multi-tone signal can be generated in which the amplitudes of the sine waves exactly follow the 50 µs or 75 µs preemphasis curve, as required. The standard FASTTEST frequency response measurement at the output of a deemphasized receiver will then be nominally flat.

Broadcast Transmission Testing

Broadcast transmission testing has several unusual aspects which require differences in audio testing techniques compared to standard methods used for amplifiers. The most obvious difference lies in the fact that the test signal injection point and the output point for measurements are often geographically separated, possibly by thousands of kilometers. Broadcasting transmission circuits often contain compressors and limiters, perhaps many cascaded compressors or sophisticated multi-band modulation level enhancement devices. The FM and TV aural transmitters contain preemphasis circuitry, with matching deemphasis circuits in the receivers. Finally, these and other problems are compounded by the fact that many broadcasting operations now operate 24 hours per day, with the programming authorities unwilling to give up any significant time periods to engineering for performance tests. Thus, routine testing of transmission quality is desired during intervals as short as one second.

Multiplex Broadcasting Measurements

Stereo signals are normally delivered to the broadcast listener by some type of multiplexing system which combines the additional necessary information for stereo reproduction onto the same carrier signal which delivers monaural audio to non-stereo-equipped listeners. The multiplex system used in high-fidelity stereo FM broadcasting is rather well standardized around the world. Broadcasting of stereo audio signals in connection with television is much less well standardized. AM (medium wave) stereo broadcasting is not widely done except in the U.S.A. FM stereo broadcasting will be briefly discussed here from a measurement standpoint; other stereo multiplexed signals are measured in a generally similar fashion.

In order to maintain compatibility with existing monaural FM receivers when stereo FM broadcasting was introduced, the primary ("baseband") signal was essentially unchanged from previous monaural practices. This was accomplished by combining the left and right audio channels in phase into an "L+R" signal and using this signal as the main frequency-modulating source of the carrier. Additionally, the left and right signals were combined out of phase ("L-R") and this signal used to modulate a sub-carrier located so that the L-R
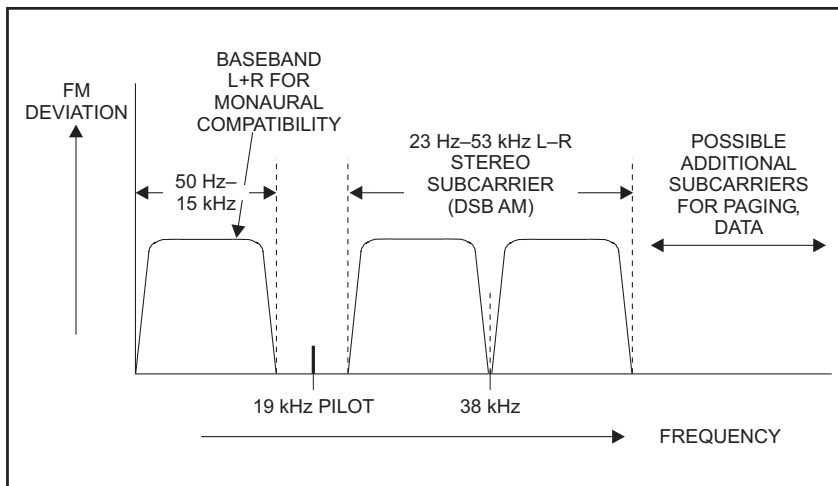
Figure 78. Spectrum, FM stereo multiplex system.

signal and its sidebands all fall above the upper range of human hearing. In the receiver, the L+R and L-R signals are processed to recover left and right output channels, which then pass through deemphasis networks. The frequency of 38 kHz is used for the sub-carrier and a double-sideband suppressed carrier AM system is used in modulation. With a 15 kHz upper modulation frequency limit for the program material, no significant energy will exist outside the range of ±15 kHz from the subcarrier, or 23 kHz to 53 kHz (see Figure 78). Low-distortion reception of a double-sideband suppressed carrier AM signal requires re-injection of the exact carrier frequency in the exact phase in the receiver. This is accomplished in the FM stereo multiplex broadcasting system by transmitting a "pilot tone" of exactly one-half the subcarrier frequency, or 19 kHz. This tone is transmitted at a low deviation level (low modulation percentage), is further rejected by low-pass filtering in the receiver, and furthermore lies above the upper frequency limit of hearing of most listeners, so it is normally inaudible. In the receiver, the pilot tone is used to regenerate the subcarrier via a phase-locked loop frequency doubler.

There are several test and measurement implications of this FM stereo multiplexing system. First, an audio-frequency spectrum analyzer (typically an FFT analyzer including coverage of frequencies up to at least 55 kHz) may be used to view the entire signal. The input signal for this spectrum analysis must be from a "composite" or "baseband" output on a demodulator, not from the deemphasized outputs. The composite or baseband output is taken immediately following the FM discriminator, and the output voltage at this point is directly proportional to the frequency deviation (percentage modulation) of the RF carrier. Thus, the baseband output amplitude can be calibrated in terms of frequency deviation and the amplitude of each signal in a display may then be directly stated in kHz deviation. The pilot tone deviation, for example, may then be compared to specifications by using the calibration factor and this type
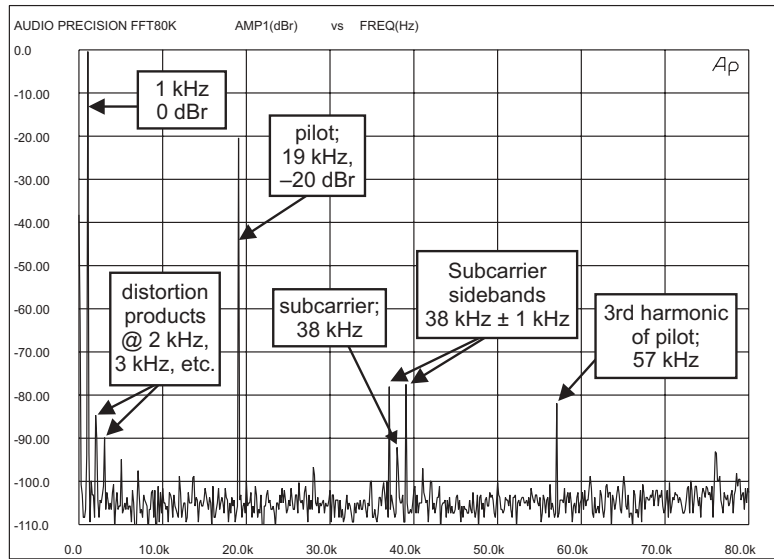
Figure 80. FFT spectrum analysis, composite FM stereo signal with 1 kHz L+R modulating signal.

Figure 79. FFT spectrum analysis, composite FM stereo signal with 1 kHz L–R modulating signal.

of spectral display. FM broadcasting transmissions may also include non-broadcast-related signals such as paging services or data communications on other subcarriers placed still higher in the spectrum where they will not interfere with the stereo subchannel.

When exactly the same signal is fed, in phase, to both the left and right in-puts of an FM stereo transmitter, the resulting modulation of the RF carrier will all take place due to the baseband signal (plus the constant low-level pilot tone) and no significant energy will occur in the subcarrier band between 23 kHz and 53 kHz. Conversely, if the same signal is fed to both left and right inputs, but with the right channel and left channel exactly 180 degrees out of phase, all deviation of the carrier will be due to energy in the subcarrier band and no energy will be found in the baseband area between 50 Hz and 15 kHz. See Figures 80 and 79 for examples of an FM stereo signal with 1 kHz modula-tion L+R (Figure 80) and L-R (Figure 79). These generator channel phase rela-tionships are used when measuring crosstalk from main-to-subcarrier and from subcarrier to main carrier, and some audio generators designed for broadcast applications provide an "L-R" output as a standard operating mode for testing sub-to-main crosstalk. Crosstalk between main and subcarriers should not be confused with stereo separation between left and right signals.

Architectures for Testing Between Separated Sites

The geographic separation between sending point and receiving point can be handled in several ways. The traditional method of manually testing such links, when long periods are available for testing, is to use two engineers who have telephone communication between them. One engineer controls the audio generator at the sending point (perhaps the studio) and one controls the ana-lyzer at the reception point (perhaps connected to an off-air demodulator at the transmitter). The engineer at the reception point directs the other engineer to set the generator to the specific frequencies and amplitudes required, then makes the measurement and records the data. To make measurements at a con-stant deviation level across the audio band, for example, the two engineers make iterative adjustments with the engineer controlling the analyzer telling the generator-controlling engineer how much and in which direction to adjust the generator. A set of measurements typical of US. broadcasting practice to completely characterize a stereo broadcasting station can take from 30 minutes to several hours by these techniques.

When automation is desired in order to speed measurements and provide better documentation, there are several possible testing architectures, each with its own characteristics, as follows:

- CCITT O.33 (formerly EBU R-27) method and other techniques relying on an FSK-coded "header" or preamble. This header is an audio fre-quency-shift-keyed (FSK) signal sent down the transmission link to identify the source, select among test sequences at the analyzer, and pro-vide an accurate synchronization signal for the start of the procedure. The header is typically followed by a sequence of single sine wave tones, usually of one second duration per tone, intended for the measurement of response, distortion, and interchannel phase. Measurement data resides at the reception point at the end of the tests. No data communication cir-

cuits are required unless the test data is to be later uploaded to another location. Time required for the standard CCITT O.33 test sequences is 31 seconds monaural or 33 seconds stereo. Other conceptually-similar techniques have reduced test sequence time to as little as five seconds by measuring at fewer frequencies and shortening each individual measurement. The technique is semi-flexible in that one of a number of previously-prepared test sequences can be chosen at the transmission point, with the corresponding previously-prepared measurement sequence selected by the FSK header. However, each sequence must first be defined, prepared, and incorporated into software at both ends of the link before it can be used.

■ A series of frequency sweeps is generated at the transmission end, combined with a series of measurements at the reception end using an adaptive method of tracking the audio signal ("external sweep"). The data resides at the reception point at the end of the tests. No data communication circuits are required unless test data is to be later uploaded to another location. Time required is on the order of 10–15 seconds for a response, phase, or crosstalk sweep, 15–20 seconds for a THD+N sweep, and 1–2 seconds for a noise measurement. The technique is inflexible in that the sequence must be defined, prepared, and installed in software at both ends of the link before it can be used.

■ Using the FASTTRIG multi-tone method with its triggering technique at the reception point which can recognize the specific multitone signal, acquire it, and perform analysis. The data resides at the reception point at the end of the tests. No data communication circuits are required unless test data is to be later uploaded to another location. Time required is approximately 0.25 s to 1 s for the burst of test signal (depending upon the low-frequency resolution desired). Several seconds of data analysis follow, during which the broadcast link has already returned to normal programming. The technique is inflexible in that the exact multi-tone signal to be transmitted must also be stored in a computer at the reception end for triggering purposes. It is semi-flexible in that the measurement sequence may be changed to add or reduce the number of measurement parameters.

■ Using a data communications link between the two ends of a single audio transmission link in order that a single engineer at either end can control the audio test instrumentation at both ends. Swept tests are plotted in real time at the control point, and real-time measurements may be observed at the control point. If remotely-controlled insertion (patch-point) switchers are installed with the audio test set at an unstaffed location and connected to the key audio circuit nodes, an engineer at the control point can make measurements or insert signal at any of those nodes. Test data is displayed and resides at the control location, which may be either the transmission or reception point, at the end of each test. Time

required is basically the same as the time required for the same tests on an amplifier or other device whose input and output are at one location; typically 10 seconds for most sweeps, 1second for a noise measurement. The technique is completely flexible; previously-prepared procedures may be run, or the operator may perform trouble-shooting in impromptu, real-time mode where he decides what his next measurement will be based on the results of the present measurement.

Emphasis

If FM or TV aural transmitters are included in the transmission circuits being tested, the preemphasis and deemphasis characteristics must be allowed for. If mid-band deviation levels are set no higher than 10% of maximum rated deviation and if the off-the-air demodulator contains deemphasis, an overall flat response should result. However, this results in the transmission circuits not being tested at anything close to their normal or peak operating levels, so the test is not rigorous. A more rigorous test is to use a deemphasized sweep from the audio generator, with levels set to maintain 100% rated deviation at all audio frequencies. The resulting frequency response from a deemphasized demodulator will then not be flat, but will show the deemphasis characteristic. A post-measurement correction with the theoretical deemphasis curve can be made to produce a nominally flat graph.

The US government broadcasting regulatory agency, the Federal Communications Commission, has always required that transmitter testing be performed at constant modulation or deviation percentages at all frequencies. The parameter then measured and plotted for frequency response testing is the output amplitude of the audio generator required at each frequency to maintain that specified value of modulation. This is actually a modulation sensitivity test. The graph is typically drawn with the vertical axis inverted so that a preemphasized transmitter will show a rise at high frequencies, rather than the actual measured fall-off in generator amplitude necessary to maintain constant deviation.

This style of testing can also be automated if the fourth architecture above is in use. The audio analyzer monitors the baseband (non-deemphasized) output of the demodulator for purposes of deviation measurement, in addition to measuring distortion and noise at the deemphasized outputs. At the baseband demodulator output, ac voltage is directly proportional to FM deviation. Software can monitor the demodulator output level and (via the data communications link, assuming that the audio generator is not at the transmitter-demodulator location) adjust the audio generator output level at each step of the frequency sweep for constant demodulator output. The result is constant deviation, and the final generator output level at each frequency is the parameter graphed for frequency response measurements.

Typical Performance, AM Transmitter & Demodulator	
Frequency response 50 Hz–10 kHz	±1 dB to ±3 dB
THD+N @ 85% modulation (22 kHz BW) 50 Hz–5 kHz	1% to 5%
S/N ratio @ 100% modulation	60 dB

Typical Performance, FM Stereo Transmitter & Demodulator	
Frequency response 50 Hz–15 kHz	±0.2 dB to ±1 dB
THD+N @ 100% deviation (22 kHz BW) 50 Hz–15 kHz	<0.05% to <0.3%
S/N Ratio @ 100% deviation	60 to 80 dB

Digital Audio and Converter Measurements

Figure 81 is a simplified block diagram of a generic digital audio device (monaural). Not all these functions are necessarily present in all digital audio devices. In the diagram, an analog input signal is converted to the digital domain via the blocks of a low-pass (anti-alias) filter, sample-and-hold circuit, and A/D converter. The digital signal may then be processed or recorded. If the device also has digital inputs, the signal from those inputs has its format converted to the digital format necessary for the processing or recording. Following the processing, or at playback in a digital recording system, the digital signal may be converted in format to the format required for a digital output. If an analog signal is required, the signal is fed to a D/A converter and low-pass (reconstruction) filter to the analog output.
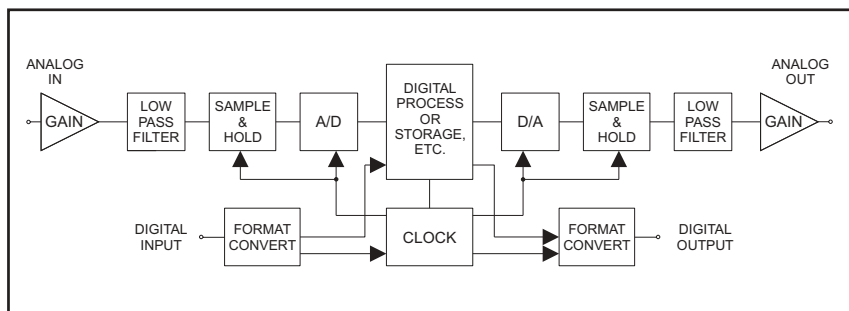
Figure 81. Generic block diagram, digital audio equipment.

Some of the tests performed in digital or mixed domains are conceptually familiar. Frequency response and distortion versus frequency are examples. Other tests may sometimes have been done in the analog domain but become

more important in the digital domain. For example, the performance shortcomings of both A/D and D/A converters are typically with very small signals rather than signals near full scale of the converter. Input-output linearity tests are good methods to measure low-level linearity of converters.

Other types of tests have been newly developed for digital devices, since "classical" analog domain tests may not be useful. For example, noise testing of analog devices is commonly done by removing the input signal and measuring the resulting output. A converter, however, may have no change in output at all when the input signal is removed, simply sending the same digital output value at each sample. In order to make realistic measurements of noise, the converter must be operating. A measurement of the quantization noise in a converter is commonly made by feeding it a full-scale, low frequency signal. At the digital output, a high-pass filter or notch followed by high-pass filter removes the input signal and its low order harmonics, but passes the remainder of the audio band for a measurement of quantization noise.

Dither

Dither is the name given to noise deliberately added to the signal before an A/D or D/A conversion in order to reduce distortion, improve linearity, and extend the dynamic range below the theoretical floor with undithered signals. The intrinsic limitations of digital conversion without dither are illustrated in Figure 82. The illustration is for an A/D converter, but the same concepts apply to D/A conversion. The converter transfer characteristic is a staircase. As the signal amplitude moves up and down, at each sampling point the instantaneous amplitude is converted to the digital output word corresponding to the nearest step of the staircase. The illustration shows only the few steps near zero. When the signal amplitude is large, the height of the individual steps is a small fraction of the signal amplitude and little error or distortion results. As
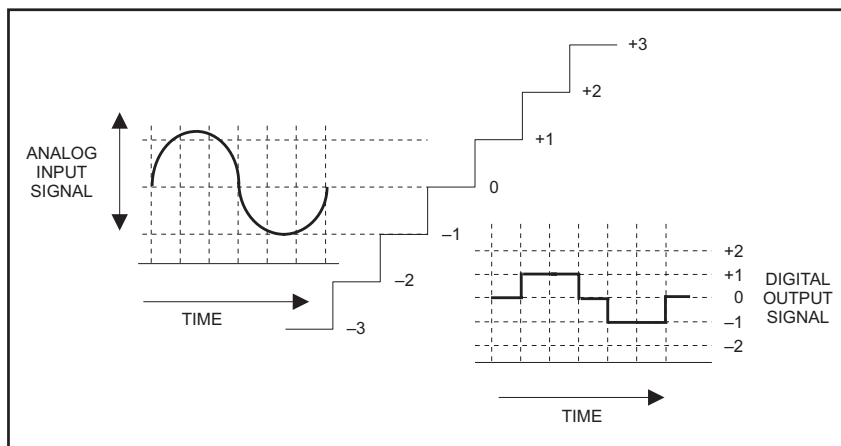
Figure 82. Transfer function, A/D converter near zero amplitude.

Figure 83. 16-bit digital domain waveform, undithered, at minimum amplitude.

Figure 84. Same signal as Figure 83, with dither.

the signal amplitude is faded down to the level where it swings across only three steps, the resulting digitized output becomes a very crude approximation of the input signal; high distortion results (see Figure 83). If the signal ampli-

tude is further faded to the level where its peak-to-peak swing is less than the height of one step, the output will not change at all and the stream of digital words at the sample rate are all identical ("digital dc at zero amplitude").

If a random noise signal of at least one step peak-to-peak is combined with the signal, all of these problems are reduced or eliminated. Now the signal amplitude can be faded below one step peak-to-peak and the digital output will still contain evidence of the signal, since the dither amplitude is enough to keep the converter digital output always changing among the three possible digital words corresponding to zero, plus one step, and minus one step. Even though the signal amplitude is still lower than the dither, it influences the composite output of signal plus dither and can be measured (with selectivity) or heard by a listener. The digital domain waveshape is no longer the three-level "square" wave of Figure 83. An instantaneous "snapshot" of a low amplitude dithered signal might look like Figure 84, with additional transitions between levels caused by the dither. The digital levels tend toward the higher value digital words when the signal is positive, and tend toward lower value words when the signal is negative. When many cycles of this noisy three-level signal are averaged, which is what the human hearing mechanism effectively does, the result is a noisy sine wave. Dither raises the noise floor of the system, but in return provides an extended effective dynamic range and lower distortion.

A/D Converter Testing

Analog to digital converters must be tested by feeding an analog signal to their input and measuring the digital output. In the early days of digital audio when no test instruments were available which measured in the digital domain, the common technique was to connect the digital output of the A/D converter to the digital input of a better D/A converter. The analog output of the D/A converter could then be measured with conventional analog test instruments. A key problem of this technique is locating the "better" D/A converter. If the A/D converter to be tested is 12 to 14 bit resolution, then 16 bit D/A converters are readily available which satisfy the "better" criterion. But if the A/D converter to be tested is at the state of the art in resolution, no better D/A converter may be available. With digital domain analyzers commercially available, finding the "better" converter is no longer necessary.

The standard battery of frequency response, THD+N versus frequency, and THD+N versus amplitude are all relevant to A/D converter testing even though the analysis is done real time in the digital domain. Response measurements in the digital domain require only DSP-implemented level measurements (rms detectors). Level measurements in the digital domain are commonly expressed in dBFS (dB below digital Full Scale) or in percent of full scale. THD+N testing in the digital domain is done with a DSP-implemented notch filter steered to track the analog signal generator frequency. THD+N may be expressed in percentage of the signal level or dB below the signal level as is customary in analog domain THD+N measurements, or it may be ex-

Figure 85. Input/output linearity, A/D converter.

Figure 86. Deviation from perfect linearity, A/D converter. (Data from Figure 85, processed.)

pressed in absolute digital domain units such as dBFS or percent of FS. Particularly for THD+N versus amplitude sweeps, using absolute units such as dBFS makes the noise-limited nature of the measurements (in a well-designed con-

verter) more obvious. A well-designed converter will exhibit essentially a flat curve of distortion when expressed in dBFS, since the quantization noise and distortion or the dither added before the conversion are constant values, independent of signal level. If the THD+N is graphed relative to the signal level, it appears to rise steadily as signal level is reduced since the constant noise becomes a larger fraction of the smaller signal level. This tends to obscure the basic noise mechanism.

An important test with converters which is not commonly made on analog domain devices is output versus input linearity. The input analog signal amplitude, at a fixed frequency, is swept (faded) from full scale down to very low levels, at or below the theoretical quantization noise and distortion level of the converter. A 16-bit linear digital audio device, for example, has a theoretical quantization noise and distortion floor of –98.03 dBFS. Therefore, the input signal should be faded through at least a 100 dB and preferably 110 to 120 dB range. The digital output amplitude is plotted on the vertical axis versus analog input signal amplitude on the horizontal axis. Whenever very low amplitude single-frequency signals are to be measured (in digital or analog domain), a bandpass filter should be used to reject as much wideband noise as possible while passing the signal without attenuation. In the digital domain, a digitally-implemented narrow bandpass filter, fixed at the analog generator frequency, should precede the DSP detectors. Figure 85 is an example of an A/D converter linearity measurement across a 120 dB dynamic range. With such a wide dynamic range graphed, it is difficult to see small deviations from linearity on the order of 1 or 2 dB or less. Post-processing of the data is useful in this case. If a straight line is fitted to a presumed-linear section of the data and then every data point of the complete set is subtracted from this straight line, the result will be a graph of deviation from perfect linearity. This deviation data can then be re-graphed on an expanded, zero-center graph as in Figure 86 to show small deviations from linearity in excellent detail. The portion of data to which the straight line should be fitted should normally be at the high level portion of the test data, but avoiding the top few dB where possible overload or compression effects could be present.

One unique attribute of digital converter testing is that noise (at least when dither is not present) cannot be measured simply by removing the signal, as with linear amplifiers. If the signal is removed and no dither is present, the output of the converter is a series of same-valued words. To measure the true quantization noise and distortion of the converter, it must be quantizing; that is, it must be converting a signal. Quantization noise and distortion is most commonly measured by using a full-scale low frequency signal, typically 50 Hz. This signal continuously exercises the converter, with the round-off errors from instantaneous input values to the nearest available digital output value producing distortion and noise across the spectrum. To measure quantization distortion in the digital domain, a DSP-implemented steep cutoff high pass filter at typically 400 Hz is used. This filter attenuates the 50 Hz fun-

damental and its low order harmonics which might typically be produced in the analog input sections, but passes without attenuation the quantization noise and distortion spread across the remainder of the spectrum above 400 Hz. An amplitude measurement of this noise level in dBFS is a measurement of quantization noise and distortion.

Noise Modulation

Noise modulation occurs in any device if the wideband noise floor is a function of the signal level. In A/D converters, low-level nonlinearity can give rise to noise modulation. Noise modulation is likely to be audible in program material such as the decay of a soft solo piano note; as it fades into silence, the noise floor may be heard to "pump" up and down.

AUDIO PRECISION   A/D CONVERTER NOISE MODULATION -70 TO -90 dB

Figure 87. Noise modulation vs frequency, digital output of high-quality A/D converter at several low-signal amplitude values.

One measurement technique for noise modulation is to make multiple spectrum analyses of noise from mid to high frequencies, each at different low amplitude values of a low frequency signal. Figure 87 was obtained by sweeping a DSP-implemented bandpass filter from 500 Hz to 20 kHz at each of five different fixed levels of a 50 Hz sine wave from 70 dB to 90 dB below full scale. If no noise modulation is present, all five spectrum analyses of noise should be identical. If the presence of signal at different amplitudes affects the noise level, this will show up as vertical offsets between the noise curves at some frequencies. The result shown in the figure is good performance; 5–10 dB vertical offsets between sweeps are sometimes found.

FFT Analysis

FFT spectrum analysis can be performed on the digital output of A/D converters with a variety of signals. Figure 88 shows the spectrum of a full-scale 1 kHz sine wave. Figure 89 shows the spectrum of a very low amplitude (–90 dB) signal, with proper dither. A number of non-harmonic signals are visible in this example.

Figure 88. FFT spectrum analysis of digital signal from A/D converter with 1 kHz full scale signal input.

Figure 89. FFT spectrum analysis of digital signal from A/D converter with 1 kHz low-amplitude signal input.

Typical Performance, 16-Bit A/D Converters	
Frequency response 20 Hz–20 kHz	±0.02 dB to ±1 dB
THD+N @ full scale (22 kHz BW) 20 Hz–20 kHz	<0.003% (–90 dB) to <0.01% (–80 dB)
Dynamic Range	90 dB to 95 dB
Linearity @ –70 dBFS to –100 dBFS	±0.3 dB to ±2 dB

D/A Converter Testing

The common frequency response, THD+N versus amplitude, and THD+N versus frequency tests are performed in the same conceptual fashion on D/A converters, except for the domains in which signals must be generated and measured. A D/A converter must be driven with a digital-domain signal, while the measurements of its analog outputs are made with conventional analog domain techniques. For the THD+N versus amplitude test, it may be more meaningful to measure and plot the results in an absolute unit such as dB relative to full-scale output, to show that the principal distortion and noise mechanisms are really independent of signal level.

Input-output linearity testing of D/A converters is done by the same concept as A/D converters. For linearity tests or any other tests at low digital signal am-

Figure 90. Input/output linearity, properly dithered D/A converter.

AUDIO PRECISION   D/A CONVERTER LINEARITY, UNDITHERED

Figure 91. Input/output linearity, undithered D/A converter.

plitudes, it is important that the digital signal generator be able to generate dither at the appropriate level. Figure 90 is a linearity sweep of a 16-bit D/A converter with proper dither. Figure 91 is the same converter, with no dither added to the digital generator output. Note that below about –90 dB with no dither, where the peak-to-peak digital signal amplitude is no longer large enough to round off to three different values of a 16-bit word, the analog output amplitude no longer follows the signal amplitude. With dither, the converter is seen to be linear for another 20 dB to 30 dB until the output signal is finally lost in the noise since the analog bandpass filter used in the measurements is not infinitely narrow.

FFT Analysis

FFT spectrum analyses can of course also be made of analog signals such as the output of a D/A converter. All analog-input FFT analyzers must convert analog signals to the digital domain via A/D converters before performing the FFT analysis, which is intrinsically a digital process. Most commercially available FFT analyzers have A/D converters of 12 to 16 bit resolution. They are therefore not useful (without assistance by pre-processing) for making spectrum analyses of the output of a 16-bit D/A converter, since the converters in the FFT instrument would add as much or more distortion as the converter being tested. If the FFT analyzer has the ability to send signal through an analog notch (band reject) filter before the signal is converted into the digital domain for FFT analysis, greatly improved performance can be obtained. Figure 92 is an FFT of a low-distortion analog signal, with many distortion products visible

Figure 92. Low-distortion analog signal with no signal processing before 16-bit converter of FFT analyzer.

Figure 93. Low-distortion analog signal with analog notch filter before 16-bit converter of FFT analyzer.

in the –90 to –110 dB range, where the 16-bit A/D converter in the analyzer is likely to add distortion. It would not be possible from this information to tell which distortion products are generated in the A/D converter and which are part of the original analog signal being analyzed. Figure 93 is an FFT of the

AUDIO PRECISION   D/A CONVERTER AT -90 dB, SPECTRUM WITH PROPER DITHER

Figure 94. FFT spectrum analysis, analog output of 16-bit D/A converter with −90 dBFS signal properly dithered.

AUDIO PRECISION   D/A CONVERTER AT -90 dB, SPECTRUM WITH NO DITHER

Figure 95. FFT spectrum analysis, analog output of 16-bit D/A converter with −90 dBFS signal, no dither.

same signal, but with an analog band reject filter inserted before the A/D converter. This notch filter attenuates the signal by about 125 dB. The remaining signal and distortion products are then amplified before being fed to the A/D converter, so that what was the -60 dB point on the original signal now be-

comes full scale on the A/D converter in the FFT analyzer. Any distortion products generated in the converter from 90 to 110 dB below converter full scale are now 150 to 170 dB below the original signal amplitude, below the noise level. Distortion products are now visible 130 to 150 dB below the original signal amplitude.

Figure 94 is an FFT of the analog output signal of a D/A converter when the converter is driven with a digital signal at 90 dBFS, the minimum undithered signal level possible in a 16-bit system. Proper dither was combined with the digital sine wave for this test. The only undesired signal visible is a –120 dB non-harmonic signal at about 2.6 kHz, with the noise level due to dither at about –128 dB. Figure 95 is under identical conditions but without dither. The noise level is now noticeably much lower, at –137 dB. But, distortion products are now visible at every harmonic through the Nyquist limit of the converter. Some of these distortion products are only about 14 dB below the –90 dB signal. Again, the distortion-reduction properties of proper dither are clearly shown.

Noise Modulation

Noise modulation of a D/A converter is tested similarly to an A/D converter as described above, except that the signal is generated digitally and the analyzer and swept bandpass filter are analog. Figure 96 is a noise modulation test of a D/A converter and shows very good performance.

Figure 96. Noise modulation vs frequency, D/A converter.

Compact Disc Player Testing

CD player testing is a special case of D/A converter testing, since D/A converters are the most critical elements in CD players. The most unusual aspect of CD player testing is that a CD player has no real-time digital input, but must be tested by playing test CDs with appropriate signals recorded on them.

Frequency response and THD+N versus frequency testing at maximum amplitude (0 dBFS) are possible with almost any test CD, since virtually all of them contain a number of tracks of full-scale signals at a variety of frequencies. The analyzer used must have the capability of making "External Frequency" sweeps, where the analyzer measures the frequency of the signal coming from the DUT, plots that varying frequency along the horizontal axis, and uses the signal frequency information to steer the notch filter for THD+N measurements.

Some test CDs also contain a series of tracks at a fixed frequency and progressively lower amplitudes. These can be used for THD+N versus amplitude and for input-output linearity testing. The analyzer must have an "External Level" capability where the analyzer measures the level from the DUT to drive the horizontal axis. However, plotting measured level vertically versus measured level horizontally does not produce a measurement of input-output linearity. The horizontal axis must instead be the original, exact levels recorded onto the test CD. There are two ways of getting this information onto the horizontal axis. If the CD player has a digital output and the analyzer has digital inputs and the ability to drive its horizontal axis from the digital input, the desired linearity measurement will result. If both these conditions are not true, the test data can still be collected by an "External Level" sweep on both axes. Then, manually or automatically, the measured values on the horizontal axis can be replaced with the exact theoretical values supplied as part of the test CD documentation.

High quality CD players often deliver frequency response flat within 0.1 dB from 20 Hz to 20 kHz. THD+N below –90 dB is available from CD players in almost any price class, and very good CD players may have distortion below –95 or –96 dB. Linearity error within 1 dB down through the –100 dB signal level is possible with the best CD players. RDAT machines can produce equivalent results when playing digitally-mastered test tapes. For combined record-playback performance, an RDAT is likely to have distortion in the –85 to –90 dB area since the A/D converters typically produce more distortion than the D/A converters. RDAT frequency response can also be flat within 0.1 dB to 0.2 dB.

Dynamic Range Testing

Another test method designed to measure quantization noise and distortion is called the dynamic range test. This test is conceptually simple; it is a

THD+N measurement (expressed in decibels) made on a signal 60 dB below full scale. The 60 dB value is then used to correct the measured value in order to produce a number describing the complete dynamic range of the instrument. For example, if the THD+N measurement at the –60 dB level is –32 dB, the dynamic range is 92 dB. The signal at –60 dB assures that the digital converter continues to operate, generating quantization noise and distortion. In practice, the dynamic range measurement may require an additional amplifier prior to most audio analyzers since the –60 dB signal level is quite low in absolute terms. On consumer CD players and RDAT machines, for example, the maximum output level with a full-scale signal is typically 2 V rms. The –60 dB level signal will thus produce 2 mV of output signal. Most audio analyzers will not make THD+N measurements on a 2 mV signal, so an external amplifier with relatively low noise and distortion will be required to bring the 2 mV signal up to the typical 25 to 100 mV level required by most analyzers. Some sophisticated analyzers are capable of measuring THD+N on arbitrarily-low signal amplitudes, even down into the tens of microvolts, and thus can measure dynamic range directly. On 16-bit devices such as CD players and the D/A sections of RDAT machines, dynamic range values of at least 90 dB are to be expected and very good devices may approach 98 dB.

Typical Performance, 16-Bit D/A Converter or CD Player	
Frequency Response 20 Hz–20 kHz	±0.02 dB to ±1 dB
THD+N @ full scale (22 kHz BW) 20 Hz–20 kHz	<0.0015% (–96 dB) to <0.01% (–80 dB)
Dynamic Range	90 dB to 95 dB
Linearity Error @ –70 dBFS to –100 dBFS	±0.3 dB to ±2 dB

Close-Coupled Acoustical Devices

Telephone handsets, headsets, hearing aids, and similar acoustical devices are typically tested with swept or stepped sine wave techniques. Handsets, headsets, and hearing aids are typically tested with very close coupling between the acoustical signal source and the microphone under test, or between the acoustical output and the measurement microphone. Often, acoustical couplers in the form of short lengths of plastic tubing are used to deliberately channel essentially all the acoustical signal from the transducer being measured into the test transducer. These couplers, or close coupling, help to reject ambient acoustical noise and to produce repeatable results test after test. No assumption is made that this close coupling actually duplicates the usage environment of the device; it is simply a standardized testing environment. The commonly-performed test is frequency response, although THD+N is occasionally measured at a few frequencies.

Telephone handset microphones are often tested with an "artificial voice." The artificial voice consists of a small loudspeaker and a monitoring microphone, plus an acoustically designed cavity such that the built-in monitoring microphone will "hear" the same sound pressure level (SPL) that the handset microphone under test will "hear." At each step of a frequency sweep, the output of the built-in monitoring microphone is used in either a software or hardware servo loop to adjust the level of the oscillator driving the small loudspeaker, so as to produce an essentially constant sound pressure level at all frequencies even though the small loudspeaker is not flat with frequency.

Hearing aids typically have compression amplifiers as part of their electronic circuits, so that the user will not be subjected to exceptionally loud sounds from very loud acoustical signals or from signals very near the hearing aid pickup microphone. The presence of a compressor in the circuit renders frequency response measurements at normal sound pressure levels useless, since the compressor works to produce constant output over a wide range of input levels. Thus, frequency response measurements are normally made at low sound pressure levels, below the threshold of compression. The monitoring microphone technique described above as part of the artificial voice permits a specific sound pressure level to be set and maintained for a sweep. Sweeps can then be run at a variety of SPLs, both below and above the compression threshold.

Microphone Testing

High-quality microphone testing is normally done in something approaching a free-field situation, similar to the usage environment of these microphones. Since a good microphone has flatter frequency response than the best loudspeakers, some technique must be used to overcome the imperfections of the acoustical source. One technique is basically the same as the artificial voice described above. A known-flat reference microphone and the microphone under test are placed side-by-side and as close together as possible, but spaced well away from the source loudspeaker and from other acoustically reflecting surfaces. A two-channel analyzer is used which has control features which automatically adjust its generator output level so as to maintain constant the measured input level on one channel. That channel is connected to the output of the known-flat reference microphone. When a response sweep is made in this mode, the result is constant sound pressure level at the reference microphone and presumably also at the microphone under test since they are very close together. If no compression is involved in the microphone under test, the same physical setup may be used without the feature of regulating sound pressure level, by measuring and graphing the difference in output of the two microphones. The sound pressure level may then vary widely due to non-flatness of the test loudspeaker. If the reference microphone is flat and the difference in output levels of the two microphones is graphed, the result will be the frequency response of the microphone under test.

The "free field" in which these microphones are tested is likely in practice to be inside a room with many acoustically-reflecting surfaces. With steady-state sine wave testing, the reflections create standing waves as reinforcement and cancellation take place. At frequencies of cancellation, the resulting sound pressure level may be near or below the room ambient acoustical noise and poor data may result. Therefore, it can be advantageous to use one of the quasi-anechoic testing methods such as MLS, impulse testing, or TDS when testing microphones.

Section 4

Audio Measurement Glossary

The definitions in this glossary are limited to the common audio industry usages of the terms. Other specialized areas may have differing definitions of some of these terms.

A-weighting filter: a specific noise-weighting filter (ANSI S1.4, IEC Recommendation 179) used to produce noise measurements which correlate well with human observations. See weighting filter.

A/D converter: a device for converting an analog input signal into a series of digital values representing the instantaneous amplitude of the signal at regular sampling intervals.

ac mains: the utility-provided ac power source from which most electronics equipment is operated. The term ac mains is most commonly used in "British" English, with "power line," "ac line," "line voltage" being common "American" English terms.

AC-2: A low-bit-rate transmission technique developed for broadcast and other professional applications by Dolby Laboratories. AC-2 is a transform-based technology.

Acceptance test limit: the poorest measured performance specification which a manufacturer is willing to accept and ship in a product.

Accuracy: the degree of conformance of a test instrument to absolute standards, usually expressed as a percentage of reading or a percentage of measurement range (full scale).

Acoustical testing: audio-frequency tests in which the signal exists as acoustical waves which must be transformed into electrical signals by a transducer such as a microphone before measurement.

AES/EBU interface: a digital interface standard for professional audio equipment interconnection.

AES: Audio Engineering Society, with headquarters in New York City.

AGC: see automatic gain control.

Alias: a nonlinear signal product falling within the audio band, caused by the presence of an out-of-band (above 1/2 the sampling rate) signal during the original A/D conversion. Alias signals fall at sum and difference frequencies of the original signal frequency (or its harmonics) and the sampling rate (or its harmonics).

Amplitude: the magnitude of a signal, which may be expressed in a wide variety of units such as volts, dBm, dBu, watts, etc. Usually, though incorrectly, considered a synonym of level.

Analog recording, processing: techniques in which a signal is represented, recorded, processed, or transmitted as a continuously variable quantity.

Anechoic chamber: a room, or space treated with sound-absorbent material so as to have no significant acoustical reflections.

ANSI: American National Standards Institute, a U.S.-based standards organization.

Anti-alias filter: a low pass filter preceding an A/D converter to prevent signals at frequencies greater than one-half the sampling rate from reaching the converter. Signals above one-half the sampling rate cannot be unambiguously converted and would appear in the digital output as signals at incorrect frequencies.

APT-X™: a low-bit-rate coding system for professional applications, manufactured by Audio Processing Technology. It divides the audio range into four bands and digitally compresses each.

ASCII: American Standard Code for Information Interchange, a standard for representing 127 alphabetic, numeric, and other characters and symbols as 7-bit binary numbers.

ATRAC: Adaptive Transform Acoustic Coding, the perceptually-based low-bit-rate coding technique used in the Sony MiniDisc system.

Attack time: the time interval required following a sudden increase in input level to a device for the output of the device to reach some stated percentage of the eventual signal level.

Attenuator: a passive network of resistors (plus possibly compensating capacitors) which reduces the amplitude of a signal by a precise amount. Often also referred to as a pad.

Automatic gain control: a technique for automatically varying the gain of a device, usually so as to maintain a relatively constant output level.

Automatic volume control: a technique for maintaining output audio level (volume) relatively constant in spite of varying input levels.

AVC: see automatic volume control.

Average responding: a form of ac signal detection indicative of the average absolute value of a waveform, typical of analog meter movements.

Back termination: a specified value of resistance or impedance connected for test purposes to a device input instead of the normal source which drives the device under test.

Balance: the degree of matching of the two sides of a balanced input or output configuration, or of the levels on the two channels of a stereo system.

Balanced line: an audio transmission line where the signal is applied differentially between two conductors, each of which has equal impedances to ground or common.

Band reject filter: a filter which attenuates the frequencies within a specified frequency range while passing essentially without attenuation the frequencies below and above that range. If the rejection range is narrow, a band reject filter is also often called a notch filter.

Bandpass filter: a filter which passes a specific frequency band essentially without attenuation while attenuating frequencies both below and above the specified band.

Bessel filters: a family of filters possessing maximally flat time delay as a function of frequency. Bessel filters are thus a time domain optimization rather than a frequency domain optimization.

Bessel function: a mathematical function which describes the amplitudes of the carrier and sidebands resulting from a frequency modulation process.

Bias: a high-frequency ac signal (typically 60 kHz or higher) combined with the audio signal in the recording head of an analog tape recorder. Proper bias level reduces the distortion of the magnetic recording process. A signal derived from the same oscillator is normally also used in the erase head to remove previous recordings from the tape.

Bin (FFT): the basic frequency resolution-determining division of an FFT-computed spectrum. The FFT process cannot resolve signals of different frequency falling into the same bin. The bin width of an FFT can be computed from (sample rate)/(number of amplitude samples) in the digitized waveform used as input to the FFT process. Bin and line are used interchangeably.

Bit: binary digit, which may have only two possible states (ON/OFF, HIGH/LOW, 1/0, etc.)

Bits of resolution: the number of bits of the binary word by which signals are represented in a digital recording or transmission system. Each bit adds approximately 6 dB to the theoretical dynamic range available. Thus, a 16-bit digital system is capable of approximately 96 dB dynamic range, etc.

Bode plot: a graph of the input-output gain and phase relationships versus frequency of an amplifier.

Bridging: a relatively high impedance input (usually balanced) which may be connected across a lower impedance circuit without significantly affecting the levels in the lower impedance circuit.

BTSC: Broadcast Television Systems Committee of the Electronic Industries Association, the group which developed the standards for the stereo audio transmission techniques used in North American television broadcasting. The technique itself is commonly referred to as BTSC stereo.

Butterworth filters: a family of filters possessing maximally-flat amplitude response as a function of frequency with no ripples in the passband. Butterworth filters can be high-pass or low-pass.

Byte: an 8-bit digital word.

C-message weighting filter: a weighting filter commonly used for noise measurements in the telephone industry. The shape of the C-message weighting filter is based on both the typical human ear frequency response and also on the response of a typical telephone receiver. The C-message weighting filter is described in the IEEE-743 and Bell 41009 specifications. See weighting filter.

Cart machine: a continuous-loop tape machine commonly used in broadcast studios.

CCIR 468: a CCIR specification which includes, among other things, standards for weighted and unweighted noise measurements. The weighted standard specifies the CCIR weighting filter and a quasi-peak detector (see weighting filter). The unweighted standard specifies a 22 Hz to 22 kHz bandwidth limiting filter and an rms detector.

CCIR-ARM: a noise measurement technique developed by Dolby Laboratories, which uses the weighting filter shape specified in CCIR Recommendation 468 but with the unity-gain frequency at 2 kHz rather than 1 kHz, and an average-responding detector.

CCIR: Comite Consultatif International des Radiocommunications (International Radio Consultative Committee), an international standards-setting organization with headquarters in Geneva, Switzerland.

CCITT O.33: a standard, adopted from the EBU R-27 recommendation, for testing broadcast transmission links. The O.33 standard starts with an audio fre-

quency-shift-keyed header which identifies the type of test sequence to follow and the source of the signal. The remainder of the test sequence is a series of tones of specified level and frequency, plus a quiet period for noise measurements. A compatible analyzer at the other end of the link is synchronized by the FSK header and then makes measurements from each section of the signal sequence. Each tone in the sequence is normally of one second in duration.

CCITT weighting filter: a noise-weighting filter described in CCITT documents P53A and P53B. See weighting filter.

CCITT: Comite Consultatif International Telephonique et Telegraphique (International Telegraph and Telephone Consultative Committee), an international standards organization for telephone communications related industries.

CD: compact disc, a digital recording system based on optically-encoded disks and linear 16-bit coding at a 44.1 kHz sample rate.

Chebyshev filters: a family of filters with a sharper rolloff than Butterworth and possessing passband ripple of specific amplitude.

Clipping: the action of a system in flattening and squaring off signal peaks when driven with a signal whose peak amplitude is beyond its linear signal-handling capability.

CMRR: see common mode rejection ratio.

CODEC: coder-decoder for reduced bit rate transmission or recording/re-production of digital audio. Codecs reduce the required bit rate by using the available bit rate resources for the most important portions of the signal. More complex coders analyze the signal and allocate the bits in a more optimal form from moment to moment. Some coders are perceptual coders, in which signals which will be psychoacoustically masked because they are nearby in frequency to a stronger signal will be encoded with fewer bits or not at all.

Common mode rejection ratio: the ability of a balanced (differential) input circuit to reject signals on the two conductors which are in phase with respect to common or ground.

Common mode voltage range: the maximum voltage between either side of a balanced pair and ground or common which can be tolerated while still permitting full common mode rejection.

Common mode: signals on a balanced pair which are in phase with respect to ground, as opposed to differential signals which are developed across the balanced pair (normal mode). Common mode signals are sometimes also referred to as longitudinal signals.

Compander: a signal processing amplifier which can compress or expand (or both) the dynamic range of the output signal with respect to that of the input signal.

Compressor: a signal processing amplifier whose gain automatically reduces with higher amplitude input signals, such that the dynamic signal range at its output is less than the signal dynamic range at its input.

Console: see mixer.

Crest factor: the ratio of a signal's peak amplitude to its rms amplitude.

Critical band: in pyschacoustics, the maximum bandwidth of noise which is perceived by humans to be the same loudness as a sine wave of the same power at band center.

Crossover distortion: a characteristic type of distortion produced in an amplifier's push-pull output stage if improperly biased such that only the peaks of low-level signals drive the amplifier into normal amplification ranges. A "dead band" input amplitude range may consequently exist, with signals in the "dead band" not producing output.

Crossover, loudspeaker: a filter network designed to separate the audio frequency spectrum into two or more ranges. This permits feeding the several drivers (woofer, tweeter, etc.) of a multi-driver loudspeaker system with only the frequency range of the signal for which they have been optimized.

Crosstalk: unwanted signal coupling from one channel of a multi-channel transmission or recording system to another.

D/A converter: a device which converts a stream of digital numbers, each representing the amplitude of a signal at a particular sampling time, into a corresponding analog signal.

Damping factor: the ratio of the rated load impedance of an amplifier to the output (source) impedance of the amplifier.

DASH: Digital Audio Stationary Head, a standard for the recording of digital audio signals onto magnetic tape.

DAT: see RDAT.

dB/octave: a standard means of referring to the ultimate rejection slope (attenuation versus frequency) of a band-limiting filter. Each pole of a band-limiting filter produces an ultimate rejection slope of 6 dB/octave (20 dB/decade). Thus, a 3-pole filter will have a rejection slope of 18 dB/octave (60 dB/decade).

dB: decibel, a ratio unit for expressing signal amplitudes. If the amplitudes are expressed in voltage, $dB = 20 \log_{10} (V_1/V_2)$. If the amplitudes are expressed in power, $dB = 10 \log_{10} (P_1/P_2)$.

dBFS: decibels with respect to digital full scale. The full scale amplitude (0 dBFS value) is the rms value of a sine wave whose positive peak just reaches positive full scale.

dBm: dB relative to a reference value of 1.000 milliwatts. dBm is thus a power unit and requires knowledge of power levels (voltage and current, or voltage and impedance, or current and impedance) rather than merely voltage.

dBr: relative dB; dB relative to an arbitrary reference value. The reference value must be stated for this to be a meaningful unit.

dBspl: dB sound pressure level, a standard unit in acoustical measurements. 0 dBspl corresponds to a sound pressure of 20 micropascal (20 micronewtons per meter squared, 0.0002 dynes per square centimeter). Thus, +94 dBspl corresponds to one Pascal.

dBu: dB relative to a reference of 0.7746 V.

dBV: dB relative to a reference value of 1.000 V.

dbx™: the family of noise reduction techniques developed by the dbx company.

DCC: Digital Compact Cassette, a low-bit-rate recording medium using the PASC perceptual coding technique, developed for consumer applications by Philips. DCC cassettes have the same physical dimensions as conventional analog cassettes, permitting design of machines which can play conventional cassettes in addition to recording and playing digital cassettes according to the DCC standard.

Decade: the interval between two frequencies with a ratio of exactly 10:1, such as the range from 20 Hz to 200 Hz, or from 1 kHz to 10 kHz.

Decimation: a mathematical process of reducing the clock rate and bandwidth content of a digital audio signal.

Deemphasis: signal processing, normally a high-frequency attenuation, to reduce the effects of noise generated in some previous sections of a transmission or recording system. In order to provide overall flat frequency response, preemphasis is performed prior to the noise-introducing section of the system. Common broadcasting deemphasis and preemphasis systems use single-pole filters and are referred to by the time constant of the resistor-capacitor network making up the circuit; for example, 50 μs or 75 μs deemphasis.

Detector: the precision ac to dc conversion section of a measurement instrument, located following all ac signal processing and prior to the indicating portion. Detectors are classified according to which parameter of the input ac signal the output dc value linearly follows—true rms, average, peak, etc.

Difference product: an intermodulation distortion signal at the frequency which is the difference between two applied signal frequencies. Sometimes referred to as difference tone. For example, test signals of 13 kHz and 14 kHz will produce a 1 kHz difference product when applied to a device which has asymmetrical nonlinearity.

Differential: balanced; a signal existing across two conductors, rather than between either conductor and ground or common.

Digital recording or processing: a technique in which the original signal is periodically sampled and the amplitude value at each sampling instant is converted into a number represented by a binary word.

DIM: dynamic intermodulation. A technique to measure nonlinearity of a device, designed to be particularly sensitive to distortions produced during transient conditions typical of program material. Also referred to as TIM—transient intermodulation distortion.

DIN 45403: a DIN specification for intermodulation distortion measurement.

DIN 45405: a DIN standard describing noise measurement, virtually identical to CCIR 468.

DIN 45507: a DIN specification for wow and flutter measurement, virtually identical to IEC 386.

DIN: Deutsches Institute für Normung, a German standards organization.

Dither: low amplitude noise, at approximately the one-half to one LSB range in amplitude, added prior to a signal quantization in order to reduce distortion, improve linearity, and extend the available dynamic range downwards below that of an undithered system of the same number of bits.

Dolby: Dolby Laboratories, a company manufacturing and licensing noise-reduction devices. Also, one of several specific noise-reduction techniques developed by Dolby Laboratories.

DSP: digital signal processor. A specialized microprocessor designed for highly efficient processing (filtering, FFT, etc.) of digitized analog waveforms.

DUT: a common abbreviation in the test and measurement field for "device under test."

Dynamic range: the difference, usually expressed in dB, between the highest and lowest amplitude portions of a signal, or between the highest amplitude signal which a device can linearly handle and the noise level of the device.

EBU: European Broadcasting Union. An organization with offices in Belgium and Switzerland, which (among other things) creates standards and recommendations for broadcast-related testing.

EIA: Electronics Industries Association

EIAJ: Electronics Industries Association of Japan:

Electronic audio tests: the class of audio tests in which all signals are already available as electrical signals, with no transducers required.

Elliptic filters: a family of filters possessing the sharpest possible rolloff characteristics for a given order of complexity. Elliptic filters differ from Chebyshev filters in having stopband zeros.

Emphasis: signal processing before and after a noisy medium or process with the goal of improving the overall signal-to-noise ratio without affecting frequency response. Emphasis systems consist of a preemphasis filter (usually high frequency boost) before the noisy section and a complementary deemphasis filter (therefore usually cutting high frequencies).

EQ: equalization, or an equalizer.

Equalizer: circuitry or equipment which produces a varying (usually adjustable) amplitude as a function of frequency.

Even order distortion: distortion products produced by nonlinearities mathematically described by even value exponents. These nonlinearities produce an asymmetrical shape in the output versus input transfer characteristic of the device. Examples include 2nd harmonic distortion and difference frequency intermodulation distortion.

Expander: a signal processing amplifier which has greater gain at high input amplitudes than at low inputs, producing an output dynamic range greater than the input dynamic range. Normally used as part of an overall system with earlier compressors, to restore the original dynamic range of a signal which was reduced in a compressor.

Fader: a variable attenuator in a mixing console or similar device, providing operator control of the level of a signal.

FFT: Fast Fourier Transform, a technique to compute the amplitude versus frequency and phase versus frequency information from a set of amplitude versus time samples of a signal.

Filter, band reject: see Band reject filter.

Filter, bandpass: see Bandpass filter.

Filter, high pass: see High pass filter.

Filter, low pass: see Low pass filter.

Filter, notch: see Notch filter.

Filter, weighting: see weighting filter.

Flat: constant gain or attenuation across a frequency band, unfiltered.

Fletcher-Munson: researchers in the early history of audio perception, who measured and documented curves of typical frequency response of the human hearing system at various sound pressure levels.

Flutter: variations in frequency of an analog-recorded and reproduced sig-nal due to short-term variations in speed of the recording and/or playback mechanism. Flutter differs from wow in that flutter consists of speed variations at a relatively rapid rate (perhaps above 6 Hz), while wow describes variations at lower rates. Wow and flutter specifications include IEC 386, DIN 45507, CCIR 409-3, NAB, ANSI C16.5, and JIS 5551.

Frequency domain: a means of representing a signal as a plot of amplitude (normally on the vertical axis) versus frequency (normally on the horizontal axis). Spectrum analyzers represent signals in the frequency domain.

Frequency shift keying: a means of transmitting digital data by shifting the frequency of a carrier signal between two values representing digital one and digital zero.

FSK: see Frequency shift keying.

Fundamental rejection: the amount, usually expressed in dB, by which a THD+N analyzer rejects the fundamental component of the input signal. The lowest measurable distortion of a THD+N analyzer is limited by fundamental rejection, along with several other attributes.

Fundamental: the lowest frequency component (normally also the highest amplitude) of a periodic signal.

Gain set: a calibrated audio-frequency attenuator, typically with rotary-switch-controlled attenuators in 10 dB, 1 dB, and 0.1 dB steps and a level indi-cator.

Gain, power: the ratio of output signal power to input signal power of an amplifier.

Gain, transducer: the ratio of output power level than an amplifier will de-liver to a load of specified resistance, to the power level that the amplifier's driving source will deliver to a specified resistance equal to the nominal input resistance of the amplifier.

Gain, voltage: the ratio of output signal voltage to input signal voltage of an amplifier.

Gate: a switch or other device which controls the passage of a signal, such as is used in a tone burst generator.

Gaussian distribution: a symmetrical probability distribution characteristic of typical random noise. The distribution follows the classic bell-shaped curve where 2/3 of the values are equal to or less than the rms signal value.

Go/no-go: a test in which acceptable results have been defined so that the test result states merely whether the device passes or fails to meet the stan-dards, rather than providing numeric results.

Graphic equalizer: an equalizer in which the gain at each portion of the spectrum is controlled by a separate fader. The faders are typically vertically-mounted slide controls, so that the positions of the control knobs forms an approximation to the frequency response of the equalizer.

Ground loop: an inadvertent signal path formed when interconnecting the chassis of two or more pieces of equipment, each possessing a safety ground. Ground loops can cause hum-related interference.

Group delay: the relative time delay between different spectral portions of a signal.

Harmonic: a spectral component at an exact integer multiple of a fundamental frequency.

High pass filter: a filter which passes all frequencies above a specified value essentially without attenuation, while attenuating frequencies below that value.

Hot pin: the pin of a normally-balanced connector which is selected to the "high side" (ungrounded) conductor when the connector is used to carry unbalanced signals.

Hum: interference at power mains-related frequencies. Hum directly at the mains frequency and odd harmonics is characteristic of magnetic coupling into the affected system. Hum at the second harmonic of the mains frequency is typically caused by inadequate filtering of power supplies with full-wave rectifiers.

IEC flutter: the method of measuring wow and flutter described in IEC 386. Virtually equivalent to DIN 45507.

IEC: International Electrotechnical Commission, a body responsible for preparing and publishing international standards for the electrical and electronics fields. The IEC is based in Geneva, Switzerland.

IEEE-488: an IEEE standard for interconnection of test and measurement instruments plus a computer/controller into an automated measurement system.

IEEE: Institute of Electrical and Electronics Engineers, a professional and standards-setting organization with headquarters in New York City.

IHF: Institute of High Fidelity, an organization which establishes standards for the testing of consumer audio equipment.

IMD: intermodulation distortion.

ISO-MPEG standards: standards defining three "layers" of performance for low-bit-rate audio. Layers one and two are MUSICAM.

ISO: International Organization for Standards, the largest of the many international groups for technical and industrial cooperation. The ISO is based in Geneva, Switzerland

JIS: Japanese standard for measurement of wow and flutter using an rms detector and very long time constants in the detector filter. The JIS standard typically produces lower numbers than the NAB or IEC/DIN standards for any specific tape machine or turntable. The JIS standard is used primarily for consumer equipment.

Jitter: the undesirable cycle-to-cycle variation in the period of a reference clock, such as are used in digital audio converters. Jitter can cause modulation sidebands and noise if converters operate from a jittered clock. Excessive jitter in an interface can cause digitally-interfaced equipment to malfunction.

Level: the amplitude of an audio signal, expressed in dB.

Limiter: a signal processing amplifier whose gain is sharply reduced above some critical threshold, so that the output signal will not exceed a specified value regardless of input amplitude.

Limiting: the action of a limiter in preventing output levels above a specified value. Also, the action of a conventional amplifier when driven beyond its linear range. In this latter context, "limiting" and "clipping" are sometimes used interchangeably except that clipping may imply a more abrupt characteristic than limiting.

Limits testing: testing in which measurements are compared to acceptable values so that a pass/fail decision may be made.

Line (FFT): the basic frequency resolution unit of an FFT-computed spectrum. Line spacing of an FFT is computed from (sample rate)/(number of amplitude samples) in the digitized waveform used as input to the FFT process. Line and bin are used interchangeably.

Line level: a relatively high amplitude range suitable for transmission of audio signals. Line level is typically in the 0 dBu to +8 dBu range.

Longitudinal balance: an alternative method to common mode rejection ratio for measuring the degree of balance of a transmission line.

Low pass filter: a filter which passes all frequencies below a specified frequency essentially without attenuation, while attenuating frequencies above that value.

LSB: least significant bit. The bit in a binary word representing the smallest possible value change.

Masking, frequency: the pyschoacoustic effect where a strong signal causes weaker signals nearby in frequency to be inaudible.

Masking, temporal: the psychoacoustic effect where a strong signal causes weaker signals occurring just before or just after the strong signal to be inaudible.

Matching impedances: assuring that the impedances of sources and loads connected together in a system are equal, so as to assure maximum power transfer.

Maximum length sequence (MLS): a pseudorandom noise sequence designed such that every possible bit combination occurs once during each repetition cycle. An MLS sequence has the property that if it is passed through a linear device under test and a cross correlation computed between the output and input of the device, the result is the impulse response of the device. The MLS is used in quasi-anechoic acoustical testing.

Maximum output level (MOL): the maximum output level that a device can deliver without exceeding a specified distortion value. Most commonly used with analog tape recorders, where maximum output level is typically defined as maximum level achievable without exceeding third harmonic distortion of 3% for a midband signal (usually 1 kHz).

MDAC: multiplying digital-to-analog converter, often used as digitally-controlled variable resistors and attenuators.

MiniDisc: a low-bit-rate recording medium developed for consumer applications by Sony. The MiniDisc system records on a magneto-optical disk and uses the ATRAC perceptual coding system.

Mixer: a specialized amplifier used in recording, broadcasting, sound reinforcement, and similar applications where a number of input signals must be combined with individual control of the level (and sometimes frequency response) of each input. Also commonly called console, mixing console, mixing desk.

MLS: see maximum length sequence.

MOL: see maximum output level.

MPEG: Motion Pictures Expert Group, a working group in the ISO setting standards for low-bit-rate video and audio.

MTS: Multichannel television sound, a generic term for the incorporation of more than one audio channel with television transmission.

Multi-track: a tape recorder with more than two tracks; typically 4, 8, or 16 for semi-professional work and 24, 32, or 48 for professional work. A multi-track recorder permits various instruments and voices to be recorded essentially independently of one another, often at different times and places. A final "mixdown" assembly into the final stereo product then tapes place before commercial distribution.

Multitone: testing techniques with stimulus signals consisting of more than one sine wave. Most multitone techniques use 15 or more sine waves distributed across the audio spectrum.

MUSICAM: a perceptually-based low-bit-rate encoding and decoding technique. MUSICAM complies with the ISO-MPEG standard.

MUX: multiplexer; a switch selecting among several signal paths.

NAB: National Association of Broadcasters, a professional-industrial organization of U.S. broadcasters which, among other things, establishes standards. The NAB has established a standard for measurement of wow and flutter.

NICAM: Near-Instantaneously Companded Audio Multiplex; a technique for audio bit rate reduction developed by the BBC in the early 1980s. NICAM is not a perceptually-based system, but uses block floating point techniques.

Noise gate: a signal processing device whose output is disabled (infinite attenuation) when the input signal level falls below a critical threshold, so that noise in the absence of significant program material is not passed to the output.

Noise reduction: a system, normally used with recording or transmission, which reduces overall noise by processing the signal prior to recording or transmission and again following transmission or a playback. Processing before recording or transmission is normally some combination of compression and high frequency boost. Processing after transmission or at playback is the opposite of the first processing in order to restore the initial dynamic range and frequency response, while reducing the effects of noise introduced in the recording or transmission medium.

Nonsymmetrical circuits, devices: unbalanced circuits or devices.

Notch filter: a band reject filter with a narrow rejection band, often used to eliminate the fundamental frequency for THD+N measurements or to reject a specific spectral component such as power mains hum or a feedback frequency in a public address system.

NTC: (1) negative temperature coefficient; the characteristic of the primary electrical parameter of a component decreasing with increasing temperature. (2) a resistor possessing a negative temperature coefficient, commonly used for temperature compensation.

Nyquist frequency: 1/2 the sample rate in a digital system, the frequency above which signals cannot be unambiguously coded.

Octave: the interval between a 2:1 range of frequencies, such as 400 Hz to 800 Hz or 5 kHz to 10 kHz.

Odd order distortion: distortion products produced by nonlinearities mathematically described with odd order exponents. These nonlinearities cause a

symmetrical shape of the output versus input transfer characteristic of a device. An example would be third harmonic distortion, or the 13 kHz product produced by 14 kHz and 15 kHz signals ($2*F_1 - F_2$).

Oversampling: a technique used in A/D and D/A converters where the sampling rate is many times higher than the minimum required for the bandwidth content of the signal. The fundamental advantage of oversampling is a simplification in requirements of anti-alias and reconstruction filters.

PAC: a low-bit-rate perceptual coding technique developed by AT&T.

Pad: see attenuator.

Pan pot: a variable attenuator used to control the proportions of an input signal which appear in each of the two stereo outputs. A pan pot can thus position the stereo image of a signal from left to right.

Parametric equalizer: an equalizer in which the center frequency, bandwidth (Q factor), and amount of gain or attenuation of a number of filters are all adjustable.

Parasitic oscillations: unwanted very high frequency oscillations that can occur in audio amplifiers, resulting in degraded noise and distortion performance.

PASC: Precision Adaptive Subband Coding, the perceptually-based low-bit-rate coding technique used in the Philips DCC system. It is similar to the MUSICAM system standardized by ISO.

Pass/fail: see go/no-go.

Passband: the frequency band of a filter in which signals are essentially unattenuated.

Patch panel: an arrangement of connectors plus flexible cables and plugs, permitting a variety of audio devices to be connected to one another as desired. Commonly used in broadcasting and recording studios.

PCM: see Pulse code modulation.

Peak-to-peak: the maximum amplitude difference between positive-going and negative-going peaks of a signal.

Peak: the maximum instantaneous excursion of a signal.

Perceptual coding: low-bit-rate coding of a digital signal according to an understanding of human perceptions of sound, so that the most important portion of the signal are coded with the greatest accuracy and coding capability is not wasted on portions of the signal which are inaudible.

Phantom power: a system for supplying dc power from a preamplifier or mixer microphone input to a microphone, on the same cable conductors which carry the microphone's output signal.

Pilot tone: a signal (19 kHz in the case of stereo FM broadcasting) transmitted as a phase reference for use in the receiver to demodulate the double sideband suppressed carrier stereo subcarrier.

Pink noise: noise whose spectral power distribution is such that the power per octave, per decade, or in any other equal-percentage section is the same anywhere across the spectrum. For example, pink noise has the same power in the octave between 50 Hz and 100 Hz as in the octave between 10 kHz and 20 kHz.

Pole: frequencies (in the context of complex variable theory) at which the transfer function of a device exhibits a maximum value.

Pot: potentiometer or variable attenuator to control the gain of an amplifier.

PPM: (1) (professional audio, broadcasting) peak program level meter. (2) (test and measurement) parts per million.

Preamplifier: amplifier stage or stages used to bring low-level signals, such as those from microphones and phonograph pickup cartridges, up to a higher, standard level suitable for signal routing, mixing, monitoring, etc.

Preemphasis: signal processing before recording or transmission through a noisy medium. Combines with later deemphasis (see) to produce overall improved signal-to-noise ratio with flat overall frequency response.

Processor: an amplifier designed to modify characteristics of a signal such as frequency response or dynamic range. Processors in common use in broadcasting and professional audio include equalizers, compressors, limiters, reverberation units, modulation processors, noise reduction units, noise gates, etc.

Proof of performance: a term in common use among U.S. broadcasters to describe the complete end-to-end performance tests (studio microphone connector through transmitter, antenna, and demodulator) to demonstrate proper operation of a broadcast facility.

Proof: common abbreviation for proof of performance.

Pseudorandom noise: noise whose amplitude-vs-time distribution appears to be random when examined over a short period of time, but which in fact exactly repeats a pattern of a certain duration. The spectrum of a pseudorandom noise signal has power only at the frequencies corresponding to integer multiples (harmonics) of a fundamental frequency whose period is equal to the repetition cycle duration. For example, a pseudorandom noise sequence with a two second repetition cycle will have a spectrum consisting of power at every harmonic of 0.5 Hz.

Psophometric filter: a filter whose response is based on the frequency response of the human hearing system. Most noise weighting filters are psophometric filters.

Psychoacoustics: the area of science combining acoustical stimulus and human response to that stimulus.

PTC: (1) positive temperature coefficient; the characteristic of the primary electrical parameter of a component increasing with increasing temperature. (2) a resistor possessing a positive temperature coefficient, commonly used for overload protection in loudspeakers and electronic equipment.

Pulse code modulation: a form of data transmission in which amplitude samples of an analog signal are represented by digital numbers.

Pulse width modulation: a form of data transmission in which amplitude samples of an analog signal are represented by the duty factor of a pulse train. Sometimes used in high power switching amplifiers.

PWM: see Pulse width modulation.

Q factor: a measurement of the selectivity or sharpness of a bandpass or bandreject filter. For a bandpass filter, the Q is equal to the ratio of the center frequency to the bandwidth at the –3dB points. For example, a bandpass filter with a Q of 5, tuned to a center frequency of 1,000 Hz, will have a 3 dB bandwidth of 200 Hz.

Quasi-anechoic: an electronic technique for eliminating the effect of acoustical reflections while testing in a normal space with reflections. Quasi-anechoic techniques all involve measuring signals whose arrival time corresponds to the acoustical propagation delay of the desired direct path, while attenuating or eliminating signals which arrive later due to the longer path length of the reflection. Quasi-anechoic techniques include impulse stimulus with gating, Maximum Length Sequence, and Time Delay Spectrometry.

Quasi-peak: a fast-attack, slow-decay detector circuit which approximately responds to signal peaks. Specifically, the detector response called out in CCIR Recommendation 468.

Random noise: noise whose amplitude-vs-time distribution is mathematically random and unpredictable, never repeating. The spectrum of a random noise signal is continuous with power at all frequencies, rather than power only at certain points as with pseudorandom noise.

RC: resistor-capacitor. Commonly used to describe filters and oscillators, as opposed to LC (inductor-capacitor).

RDAT: rotary head digital audio tape, a standard originally developed as a consumer medium for recording of digital audio signals but now primarily used in professional applications.

Real-time analyzer (RTA): a simple audio spectrum analyzer consisting of a number of bandpass filters tuned to staggered frequencies across the audio spectrum, each filter followed by a separate detector and indicator. Typical real-time analyzers are made up of one octave bandwidth filters (thus about 10 bands across 20 Hz–20 kHz), 1/2 octave filters (20 bands), or 1/3 octave filters (about 30-31 bands across the audio spectrum). Real-time analyzers are commonly used with microphones and a pink noise source or program material for acoustic measurements or monitoring. They lack the selectivity, resolution, and accuracy for most electronic audio device measurements.

Reconstruction filter: a low-pass filter following a D/A converter, used to remove the high frequency clock signal and smoothly integrate between the discrete voltage values from the D/A.

Record (digital sampling): a section of sampled signal in memory, represented as binary numbers.

Recovery time: the time interval required following a sudden decrease in input level to a device for the output of the device to reach some stated percentage of the eventual signal level.

Rectangular probability function dither: (see dither) dither which has equal probability of occurrence at any amplitude value between plus one-half LSB and minus one-half LSB deviation from the nominal value. Thus, a graph of probability versus digital value is a rectangle.

Residual distortion: the irreducible minimum distortion of an audio generator and analyzer. Residual distortion is thus the "floor" below which the instrument is not useful; measurements above but approaching within 6 to 10 dB of the residual distortion value are less accurate.

Residual noise: the irreducible noise in a measurement instrument which sets a floor for amplitude measurements.

Resolution: the smallest change in a measured parameter to which a measurement instrument can respond.

Return loss: a measure of power reflected back to the originating end of a channel due to impedance mismatches in the channel.

Reverb: reverberation.

Reverberation: acoustically, the effect caused by acoustical waves traveling various paths (direct path from sound source to listener location plus some number of reflections) with the reflected power being delayed in time due to greater path length. Many signal processors simulate reverberation electronically.

Reverse termination: a resistance or impedance connected at the input to a device under test in order to make repeatable noise measurements.

Ripple: (1) undesired ac variations on a dc power supply output. (2) variations in frequency response in the passband of a filter.

rms: root mean square, the preferred form of ac signal detection which measures amplitude in terms of its equivalent power content, regardless of signal waveshape.

Rolloff frequency: the frequency considered the transition between the attenuated and non-attenuated portions of the frequency response of a high-pass or low-pass filter. The filter has a specified attenuation, usually either 3.01 dB or the magnitude of the ripple, at the rolloff frequency.

RSS: root sum square; a method for combining the power of a number of signals by squaring each, summing (adding) these squared values, and finally extracting the square root of the sum.

RTA: see Real time analyzer.

Rub and buzz: a variety of distortions and noises created in a loudspeaker, mostly due to mechanical defects such as voice coil rubbing the magnet, cone touching connection wires, etc.

S/N ratio: see signal to noise ratio.

Sample clock: the time reference signal in a digital system which determines the intervals at which the signal will be sampled.

Sample frequency, sample rate: the frequency at which the signal is sampled in a digital system. The sample rate must exceed twice the highest analog frequency to be converted. Commonly-used sample rates are 48 kHz, 44.1 kHz, and 32 kHz.

SAP: secondary audio program; a third channel in the BTSC stereo system which may be used for multilingual broadcasting or other services.

Saturation: the effect occurring when driving an audio device (particularly an analog magnetic tape recorder or transformer) beyond its linear region, where the output amplitude cannot linearly follow the input. Used somewhat interchangeably with the terms "limiting" and "clipping."

SCA: an auxiliary multiplexed channel or channels in a broadcast FM signal, located higher in frequency than the stereo sub-channel.

SDIF: Sony Digital Interface, a serial digital interface using three coaxial cables. The SDIF and SDIF2 (a later version) were originally used on the Sony PCM-1600 series of digital audio processors.

Second order products: distortion products produced by a term with the exponent "2" (i.e., a squaring function) in the device transfer function.

SEDAT: a low-bit-rate coding technique developed by Scientific Atlanta.

Separation: the isolation (usually stated in dB) between the two channels of a stereo device.

Shaped dither: (see dither) dither in which the power-versus-frequency distribution has been modified to place most of the power at frequencies where the human hearing system is not as sensitive.

Shielded cable: a construction technique for cables in which a metallic outer conductor completely surrounds one or more signal-carrying conductors. This outer shield is normally connected to ground (earth) at one or both ends.

Short circuit: zero-Ohm connection.

Sigma-delta: an A/D converter technique characterized by a relatively low resolution quantizer operated at very high speeds, digitally low-pass filtered to obtain the final digital signal.

Signal-to-noise ratio: the difference in level between a reference output signal (typically at the normal or maximum operating level of the device) and the device output with no signal applied. Signal-to-noise ratio is normally stated in dB. The device input conditions for the noise measurement must be specified, such as "input short circuited" or with a specific value of resistance connected at the device input instead of a signal.

Single-ended: unbalanced.

Skirt: the portion of response curve of a bandpass, high-pass, or low-pass filter where the attenuation increases at a rapid rate.

Slew rate limiting: the condition when an amplifier output is unable to follow the rate of change of the input signal.

Slew rate: the rate of change of a signal amplitude, typically expressed in Volts per microsecond.

Slider: a particular design of attenuator or potentiometer in which the control know travels in a straight line (as opposed to rotary attenuators).

SMPTE RP120: a specification for intermodulation distortion measurements.

SMPTE: Society of Motion Picture and Television Engineers. A professional and standards-setting organization.

SPDIF: Sony Philips Digital Interface; a digital interface for consumer audio equipment. Sometimes also referred to as the EIAJ interface. The SPDIF is similar to the professional AES/EBU interface, but is normally an unbalanced coaxial signal of lower amplitude. Most of the status byte definitions are different between SPDIF and AES/EBU.

Spectrum analyzer: a measurement instrument which measures and displays a signal as a frequency domain representation of amplitude (usually verti-

cal scale) versus frequency (horizontal scale). Spectrum analyzers may be designed with a variety of technologies including real-time analyzers (bandpass filter bank), heterodyne or scanning analyzers with a local oscillator, mixer, and selectivity at a fixed intermediate frequency, scanning tunable filters, and FFT analyzers.

Split site: a broadcasting operation with studio at one location and transmitter at another. Testing of split-site systems imposes some constraints on test equipment and techniques because signal must be injected at one location while measurements are made at another.

Status bits/bytes: bits in each subframe of an AES/EBU or SPDIF/EIAJ serial audio signal which are assembled into bytes used for control and information. Signal characteristics such as signal emphasis, copy protection, sample rate, etc., are communicated in the status bytes.

Stereo generator: a device which converts left and right channel audio into a composite multiplexed signal suitable for modulating an rf carrier.

STL: see studio to transmitter link.

Stopband: the frequency band across which a filter has at least some minimum specified value of attenuation.

Studio to transmitter link: a transmission path from program material from broadcast studio to transmitter, usually a microwave link.

Sum product: an intermodulation distortion product at a frequency equal to the sum of two input frequencies.

Symmetrical: balanced.

TDS: see time delay spectrometry.

Terminating: connecting the specified load resistance or impedance to a device.

Termination: a specific resistance or impedance value which must be connected to the output or input of a device under test for certain parameters to be measured.

THD: total harmonic distortion. Normally computed from a series of selective measurements of the amplitudes of all significant individual harmonic distortion products. The bandwidth of each measurement must be sufficiently narrow that noise has no significant effect on the measurement.

THD+N: total harmonic distortion plus noise. Measured by attenuating the fundamental signal with a narrow-band notch filter, then measuring the remaining signal which consists of harmonics of various order, wide-band noise, and possibly interfering signals. This is the common harmonic distortion method implemented in most analyzers.

Third octave: a bandwidth of 1/3 octave, or a frequency ratio of 1.2599:1. Three successive frequency changes by this ratio result in a total frequency change of 2:1 (one octave).

Third order: distortion products produced by a cube (exponent of 3) term in a devices's nonlinear transfer function.

Three head: a tape recorder design including erase, recording, and play-back heads (in that order) so that a just-recorded signal may be monitored from the tape during the recording process.

TIM: transient intermodulation distortion. See DIM (dynamic intermodulation distortion).

Time delay spectrometry (TDS): A technique for quasi-anechoic frequency response measurements, in which a sine wave stimulus is linearly swept in frequency while a bandpass filter in the analyzer is swept at the same rate, but delayed in time by the acoustical propagation delay between the loudspeaker under test and the measurement microphone. As a result, the direct (first arrival) signal is passed unattenuated while reflected signals are attenuated since they arrive later (due to longer path length) when the filter has swept onwards to another frequency.

Time domain: a means of representing a signal as a graph of amplitude (usually on the vertical axis) versus time (on the horizontal axis). An oscilloscope produces a time domain representation of a signal.

Toneburst: a succession of sinewaves preceded and followed either by no signal or by the same sine wave frequency at a lower amplitude. Tonebursts are used to simulate program material during calibration of peak meters and in performance tests of audio processors, and for transient response testing of loudspeakers.

Transducer: a device to change acoustical power into an electrical signal (microphone, for example) or electrical power into an acoustical signal (loud-speaker, for example).

Triangular probability function dither: (see dither) dither in which a graph of the probability of occurrence of an amplitude rises linearly from zero at values plus-or-minus one least significant bit above or below the nominal value, to unity at the nominal value. The graph thus forms a triangle of unity value at the horizontal center (nominal digital value on the horizontal digital value scale), falling to zero at one LSB to the left or right on the horizontal scale.

Trigger: an event which causes another event or action, often initiating a signal generation or acquisition.

Tweeter: a loudspeaker unit optimized for performance at high (treble) frequencies.

Twisted pair: an audio cable consisting of two similar conductors twisted together to help assure that any noise voltages coupled into the cable will be in equal amplitude between each conductor and ground (earth), therefore representing a common mode signal which will be substantially rejected by a high quality balanced input.

Two sigma flutter: the value of wow and flutter which is exceeded by only five percent of the instantaneous readings over a period of time.

Two track: tape recorder for stereo recording and/or playback.

Unbalanced: an audio connection in which the desired signal is present as a voltage with respect to ground or common, rather than as a differential signal across a pair of balanced conductors.

VCA: see voltage controlled amplifier.

VCO: see voltage controlled oscillator.

Voltage controlled amplifier: an audio amplifier with a dc voltage control port which electronically varies the gain of the amplifier.

Voltage controlled oscillator: an oscillator whose frequency can be varied via a dc control voltage.

VU meter: a volume unit meter, used to indicate program levels in broadcasting, recording, and similar applications.

Weighting filter: a filter with varying attenuation as a function of frequency so as to produce a measurement where the various spectral components affect the measurement in a specified fashion. Most commonly-used weighting filters are attempts to correspond to the varying response of the human hearing system in order to produce measurements (usually of noise) which correlate well with human observations.

White noise: noise whose spectral power distribution is such that there is equal power per Hz, anywhere in the spectrum. For example, white noise will have the same power in the 30 Hz bandwidth between 70 Hz and 100 Hz as in the 30 Hz bandwidth between 10,000 Hz and 10,030 Hz.

Window: an amplitude-vs-time function used to multiply the corresponding samples of a digitized waveform before computing an FFT. Window functions go to zero at the two ends of the record. This provides better selectivity, reducing the signal spreading when non-windowed FFTs are computed on signals which are not exactly synchronous in the signal buffer length being transformed.

Woofer: a loudspeaker unit optimized for performance at low (bass) frequencies.

Word length: the number of bits in a word in a digital system.

Wow: a low frequency variation in pitch of an analog recorded and reproduced signal due to low rate speed variations of the mechanical system (disk or tape). See also flutter. Wow and flutter specifications include IEC 386, DIN 45507, CCIR 409-3, NAB, ANSI C16.5, and JIS 5551.

XLR: a high-quality balanced connector designed for audio applications. Sometimes also called a Cannon connector.

Zero: frequency (in the context of complex variable theory) at which the transfer function of a device has zero value.

Index