

In This Issue...

The Digital Interface	
By Julian Dunn	1
APWIN 2.14 now available	
What's new in this release	9
Meet Dave Post	
New Sales Support Engineer	10
New TECHNOTES	
TECHNOTES 28 and 30 are now available	10
Audio Events Calendar	
Audio Precision at NAB and AES	11
Turning Autofill Off	
For power users	11
Monitor Outputs Revisited	
Signal levels to expect	11
Personalities	
Mike Wolfe	12

TECHNOTE Series by Julian Dunn

Audio Precision is pleased to be publishing an extract of another in the series of TECHNOTES written by one of the world's premier designers and consultants in digital audio, Julian Dunn. Like all Audio Precision technical material, you will be able find Julian Dunn's TECHNOTES on the Audio Precision Web page at www.audioprecision.com as PDF documents. If you would prefer hard copy, please contact us at sales@audioprecision.com or 1-800-231-7350.

All the TECHNOTES in this series were condensed and printed as installments in *AUDIO.TST*, beginning with "Jitter Theory—Part 1" in the December 1999 issue. **Ap**

The Digital Interface

by Julian Dunn

This article is the first of two installments excerpted from TECHNOTE 26: The Digital Interface. The full TECHNOTE runs much longer with more topics and illustrations, fuller discussions and complete references.

The companion pieces TECHNOTE 23: Jitter Theory, TECHNOTE 24: Analog-to-Digital Converter Measurements, and TECHNOTE 25: Digital-to-Analog Converter Measurements, also by Julian Dunn, have much in common with this article and with TECHNOTE 26. You may want to refer to these documents.

Visit the AP Web site for the complete text and downloadable procedure files for TECHNOTES 23, 24, and 25. TECHNOTE 26 will be available in the late spring.

This article describes the basic digital interface as defined by the Audio Engineering Society (AES) and the International Electrotechnical Commission (IEC), and also includes synchronization topics, channel status charts and an overview of the applicable standards; the excerpt in the next issue of *Audio.TST* will look more deeply into the testing of the digital interface to verify conformance.

Introduction

The AES3 and IEC60958 standards provide a common interface for digital audio signals.

The interface defined in AES3 and IEC60958-4 is commonly called the




Figure 1. Interface waveform with and without termination.

“professional standard;” IEC60958-3 defines the “consumer standard” interface.

There are a number of differences between the professional and consumer standards which in some cases can render them completely incompatible. For proper performance, the consumer and professional interfaces should not be mixed. However, they are similar enough that in many situations, given the right electrical connections, the embedded audio can be carried from one standard to the other.

By requiring conformance with these standards, a user of digital audio equipment rightfully expects compatibility. Compatibility allows interconnecting the equipment without suffering loss of performance or functionality—which is, after all, the aim of interface standardization.

The digital audio interface carries three types of information:

- timing,
- audio data, and
- non-audio data.

Some of this information can be degraded by conforming but non-ideal implementations of the interface. We shall consider aspects of interface behavior and performance that may make one implementation more useful than another, such as the ability of a receiver to tolerate incoming jitter or a wide range of frame rates, or the precision with which a transmitter maintains synchronization.

This TECHNOTE also discusses synchronization. For real-time applications, such as recording, replay, or transmission it is important to have sample synchronization between equipment. The AES11 specification is a useful basis for defining good practice in this area, and TECHNOTE 26 describes the principle by which AES11 can produce a form of synchronization.

BASIC INTERFACE FORMAT

Bi-phase coding

Without modification, pulse code modulation (PCM) audio data can be difficult to transmit and receive reliably. Consider the case where all the bits are set to ones (or zeros) for a period of time. This is essentially DC, and to pass this signal accurately the electrical circuits associated with the interface would need to pass DC. Also, with a “digital DC” signal, there are no (or few) variations to suggest the clock rate, and a separate sync connection would be essential.

To address these interface problems, the digital data is encoded using *bi-phase coding*. This coding scheme carries the clock that is used for reproducing the sampling frequency for recovering the data. Since the mean voltage on the interface is zero, it also

eliminates DC, with the result that the data can be AC-coupled through a transformer or series capacitor. The coding works like this:

Each data bit has a *time slot* that begins with a transition and ends with a second transition, which is also the beginning transition for the next time slot. If the data bit is a “1,” an additional transition is made in the middle of the time slot; a data “0” has no additional transition.



Figure 2. Bi-phase coding.

This means that even with a digital DC signal of data zeros, there are still transitions at each time slot (or bit.) The clock is always carried by these regular transitions, the interface signal is now clearly AC, and the direction of the transitions (or signal polarity) becomes irrelevant. Figure 2 illustrates this bi-phase coding with 6 bits of data.

Unit interval

Many of the timing parameters on the interface are defined in terms of the *unit interval*, or UI. This is the shortest nominal interval between transitions. The bi-phase coding introduces a second transition (for data “1”) into the time slot, which means that a time slot is defined as 2 UI wide.

Framing

The data carried by the interface is transmitted serially. In order to identify the assorted bits of information the data stream is divided into *frames*, which are 64 time slots (or 128 UI) in length. Since the time slots correspond with the data bits, the frame is often described as being 64 bits in length, but the *preamble* sections (see below) break this correspondence.

Each frame consists of two *subframes*. Figure 3 shows an illustration of a subframe, which consists of 32 *time slots* numbered 0 to 31. A subframe is 64 UI in length.

The first four time slots of each subframe carry the preamble information. The preamble marks the subframe start and identifies the subframe type.

The next 24 time slots carry the audio sample data, which is transmitted in a 24-bit word with the least significant bit (LSB) first.



Figure 3. The AES3 subframe (24-bit audio data).

After the audio sample word there are four final time slots, which carry:

- the validity bit
- the user data bit
- the channel status bit, and
- the parity bit.

The two subframes in a frame can be used to transmit two channels of data (Channel 1 in subframe 1, Channel 2 in subframe 2) with a sample rate equal to the frame rate, or they can instead carry successive samples of the same channel of data, but at a sample rate that is twice the frame rate.

Preambles

A preamble is a distinctive data pattern carried in the first 4 time slots of a subframe to mark subframe and block starts. There are three preambles, all of which break the bi-phase coding rule by containing one or two pulses which last 3 UI. This rule-breaking means that the pattern cannot occur anywhere else in the pulse stream.

Subframe 2 always begins with a Y preamble. Subframe 1 *almost* always begins with an X preamble, with this exception: every 192 frames the X preamble in subframe 1 is replaced with a Z preamble, which indicates a *block start*. This provides framing of the information carried in the channel status fields—the channel status block.

Because the interface signal is valid with normal or inverted polarity, the preambles can be viewed as starting with a falling transition:




Figure 4. Preamble patterns with a falling 1st transition.

or as starting with a rising transition:




Figure 5. Preamble patterns with a rising 1st transition.

Under the bi-phase coding rules there should be a transition between each time slot but for each preamble there are two time slot boundaries without transitions. The first of these bi-phase coding violations is in the same place for each preamble—after time slot 0. This identifies that a new subframe has started. The following pattern then identifies which subframe it is.

The time slot numbers in Figures 4 and 5 correspond with the bit-position numbers shown in Figure 2 and are 2 UI wide. The preambles are 8 UI wide

and so take the same amount of time as 4 bits.

Audio data

After the preamble the audio data is transmitted with the LSB first. For audio word lengths less than 24 bits the data is justified to the most significant bit (MSB) and zero-filled below the LSB, as shown in Figure 6.




Figure 6. The AES3 subframe (16-bit audio data).

In some of the audio modes—those that transmit 20 bits or fewer of main audio data—the first four bits after the preamble can be used for another signal known as auxiliary audio data. This has the subframe structure of Figure 7. If this auxiliary data is used, then the channel status (see page 7) should indicate that the maximum word length is 20 bits, and the receiver should mask off the auxiliary audio field so that any values there are not added to the main audio sample values. Unfortunately, many decoders are not that sophisticated.



Figure 7. The AES3 subframe (20-bit audio data with aux data).

The use of auxiliary audio is very rare. One application is for voice communications, and AES3 suggests the auxiliary bits can be used for coordination purposes. One way of doing this is to transmit a 12-bit channel at a sample rate of one-third of the frame rate. Other applications, like the use of the auxiliary audio field for transmission of a data-compressed version of the main audio signal, may also be possible.

Validity bit

The validity bit was originally intended to somehow qualify the transmitted data. If the bit is set then the data is identified as not suitable for conversion (to analog audio). However there are some applications that set the validity bit if they have found an error and concealed it. This behavior is quite common for CD players, for example.

This confusion as to the function of this bit means that it is not easy to decide how a receiver should behave when a sample is marked as invalid.

When the IEC60958 or AES3 stream is used to transmit data that does not represent linear PCM audio, then the bit should certainly be set. This has a chance of causing linear PCM replay equipment to mute, which is preferable to an attempt to reproduce the data as an audio signal.

The specifications for carrying data-compressed audio on AES3 or IEC60958 require this bit to be set, so that linear PCM receiver devices will recognize that they need to mute. This has the potential benefit of stopping the receiver from producing a burst of high-level noise from of the data before the channel status pattern can identify the signal as not being linear PCM audio by reading the channel status *non-audio* bit.

User bit

The user bit can be used to carry user-specific information. In practice this means application-specific information for consumer devices such as CD or DCC.

The consumer specification, IEC60958-3, has defined a packet-based format for carrying program-related information in the user data stream and defines rules for the preservation of the user data by various classes of equipment.

Consumer format user data transmission classes and behavior		
Class	Equipment Behavior	Examples
I	Generating original user data	CD, DAT, DCC, mini disc
II	User data transparent (or no user data output)	Sound processor
III	Mixed mode user data (transparent but with some exceptions)	Mixer, sample rate converter, sound sampler

In the consumer format the user data streams from subframe 1 and subframe 2 are combined to form one stream at 2 bits per frame. This means that for a frame rate of 44.1 kHz the user data rate is 88200 bits/sec.

The professional specification, AES3 (and IEC60958-4) has channel status patterns that allow the indication of various formats of user data:

- 192-bit block with same block start as channel status.
- AES18 (packet based).
- IEC60958-3 user-data format.

In the professional format (except for the IEC60958-3 user-data format mode) the user data streams from subframe 1 and subframe 2 are associated with the audio channel being carried in that subframe. Therefore there are two streams, each with one bit per frame. (The case of user data transmission with single-channel double sample rate mode is not defined explicitly. It may be logical to combine the user data stream into one with 2 bits per frame, as both subframes carry the same channel.)

At the time of publication there were few applications of the professional format that use the user-data channel—so subtleties of user data implementation may not be useful.

Channel status bit

The channel status bit is transmitted in a block of 192 bits. A frame starting with preamble “Z” (see above) identifies the first bit of the block. The Z preamble is sometimes called “block start.”

There are independent channel status bits for each subframe 1 and subframe 2 so there are actually two blocks. Quite often these two blocks carry identical data, and many receivers only examine the data from one of the subframes.

Some of the channel status bits affect how equipment should treat the data in the audio sample word. In particular the “non-audio” and “emphasis” fields make significant differences to the way the data needs to be interpreted.

If the non-audio bit is set then the audio sample word is not suitable for decoding as linear PCM data. The name “non-audio” is a bit of a misnomer, as audio using data-compressed formats, such as MPEG, DTS, and Dolby AC-3, are flagged as “non-audio” because treatment of their raw data as if it was linear PCM would be inappropriate and would result in the generation of high level noise. The standard for carrying these data-compressed formats is IEC61937, for consumer applications, or SMPTE 337M, for professional applications.

If the emphasis field indicates that the signal has “emphasis” then “deemphasis” should be applied in any conversion to analog. The only emphasis supported by the consumer format is the CD type. This has a high-frequency boost shelf with time constants of 50 ms and 15 ms for the zero and pole. The professional format supports this format as well as J-17 emphasis, which has time constants at approximately 333 ms and 38.5 ms.

Apart from the first two bits, the meaning of the bits within the block is defined differently for the consumer and professional formats. More detailed information can be found in the full text of *TECHNOTE 26*.

Parity bit

The parity bit is used to maintain even parity for the data as a means of error detection. Specifically, even parity in the interface signal means that there is an even number of mid-cell transitions in the data area, which spans time slots 4 to 31. Since there is an even number of all other transitions, even parity means that there is an even number of transitions in every frame.

Even parity has the effect of starting each subframe with a transition in the same direction all the time. As a consequence, the transmitter of an AES3 or IEC60958 stream does not need to calculate parity, and the

receiver needs only to verify (since the parity bit is the last bit of the subframe) that the state of the second half of the parity bit is always the same as its state in the previous subframe.

If an error occurs, it is most likely to be a pair of missing transitions as an individual pulse was not detected. If a pair of transitions are missing, the parity will not change. In fact, in many schemes for decoding the bitstream a genuine parity error is impossible.

However, a violation of the bi-phase coding could be detected at such a point, since at least one of the missing transitions would be on the time slot boundary. It is much more useful to check bi-phase coding violations to identify errors than to use the parity bit.

Electrical properties

There are three basic electrical formats:

- The balanced format. This is the primary professional format and is defined in AES3.
- The consumer coaxial format. This is defined in IEC60958-3.
- The professional coaxial format (defined in AES-3id and in SMPTE276M). This was developed to use analog video transmission systems for digital audio transmission.

Balanced format

This uses a shielded twisted pair cable to carry the interface signal differentially and is normally coupled to equipment with a standard XLR connector. (See IEC60268-12). This has the advantage that we can use cabling that is common with analog interfaces. However, this can result in confusion between the two types of connections.

Though not required by AES3, many




Figure 8. XLR connectors for balanced format AES3 interface.

designs use small pulse transformers at the receiver and transmitter. In the same way as the balanced interfacing application for analog audio, the transformers offer advantages for reducing emissions and susceptibility to inductive coupling as a consequence of

the improved current balance in the line.

Transformers are required by the EBU version of the specification, EBU 3250. This is motivated by the need to maintain a high common-mode impedance at the cable terminations so that crosstalk is minimized. Crosstalk is of particular concern for EBU members because of the large amount of cabling run in parallel at broadcast installations.

Like the other electrical formats there is a requirement for the cable impedance and the transmitter and receiver termination impedances to be matched. In this case the nominal impedance is 110 Ω .

At the transmitter, the amplitude of the signal should be between 2 V and 7 V peak-to-peak with the output terminated. Without termination (assuming conventional implementation of source impedance) the generator voltage would be twice that, as shown in Figure 9. This can be driven from complementary outputs with logic operating from 3.3 V or 5 V rails and using a 1:1 transformer. A line driver circuit is shown here:




Figure 9. AES3 balanced format line driver.

At the receiver, the amplitude of the signal may be significantly reduced through cable losses; or, there may be no loss at all. As a consequence the range of possible receiver interface voltage levels is much greater than at the transmitter.

These losses have a greater effect on the high-frequency component of the signal, with the result that the heights of the shorter pulses will become lower than that of the longer pulses. This distorting effect on the signal means that it is not adequate to refer to the “peak-to-peak amplitude” of the signal but instead to measure the size of the “eye” in an eye diagram, which is described later.

It is possible to have an equalization circuit to compensate for some of the distortion, and this would be fitted prior




Figure 10. AES3 balanced format receiver with equalizer.

to the differential-to-single-ended data slicer, as shown in Figure 10.

Though they were popular in the early applications for AES3, equalizers are not used very often in modern designs. This may be because there is an expectation that cable losses will not be as significant, perhaps because lower-loss cable is used; and also because in most applications the cable length is quite short—significantly less than 100 meters. In the early 1980s I found that with the standard BBC-specified shielded twisted pair cable used for analog audio signal distribution, it was possible to get reliable operation over 100 meters without an equalizer, and that this could be extended to 250 meters with an equalizer. Moreover, with short cable lengths, the equalizer can be a liability, increasing the sensitivity to errors from cable reflections due to impedance mismatch.

Unbalanced

The two unbalanced formats use a 75 Ω impedance-matched coaxial cable for transmission.




Figure 11. RCA “phono” connector for consumer interface.

The consumer version has a transmitted level of 0.5 V peak-to-peak and uses the same kind of coaxial connector (the RCA “phono” connector, defined in IEC 60268-11 table IV sec.8.6) that is used for consumer analog connections.



Figure 12. BNC connector for unbalanced professional interface.

The professional version has a level of 1 V peak-to-peak and uses a BNC connector (see IEC 60169-8).

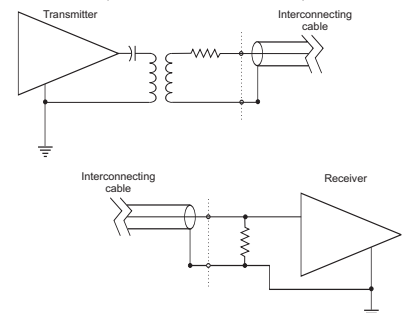


Figure 13. Unbalanced format transmitter and receiver.

The same kind of interface signal distortion occurs in the unbalanced version as for the balanced version of the interface, so the eye diagram is also used here in to assess and define signal levels and receiver characteristics.

In the consumer application (IEC60958-3), short lengths (perhaps 3 feet or 1 meter) of cable designed for analog audio interconnections will work quite adequately, even though the cable transmission characteristics are poor.

The professional specification (AES-3id) is intended for use over much longer distances, and uses professional analog video cable. 75 Ω video coax cable has an appropriate frequency characteristic for this application, and long transmission distances are possible. AES-3id illustrates that transmission distances of more than 2 km can be achieved with sophisticated equalization schemes.

Optical

In common use for the consumer format is an optical interface called TOSLink®, after the version sold by Toshiba. This uses plastic multi-mode optical fiber with a red light emitting diode transmitter and a photo diode receiver. The transmission distance is limited to less than a few yards (or meters). IEC60958-3 has a section for defining this format but it is still “under consideration.” As a result, methods of defining receiver and transmitter performance do not have a benchmark to evaluate against.

There are two connector formats for the optical fiber. The older and more widely used uses a friction lock connector type F-05 specified in IEC60874-17, shown in Figure 14.




Figure 14. TOSLink® optical connector.

This connector is too large for portable audio equipment, so a coaxial connector has been developed that appears quite similar to the electrical 3.5 mm mini-jack plug used for personal stereo headphones. This is shown in Figure 15. The socket for this connector has the advantage that it can double-up as the analog headphones jack and hence use no extra space on the equipment surface.

SYNCHRONIZATION

The embedded clock defined by the interface bit-cell transitions, the




Figure 15. 3.5 mm optical connector.

subframe and the frame boundaries can be used as a timing reference by equipment to derive timing for converters, processors, and digital outputs. For digital outputs AES11 defines limits for the timing offset between the frames of the reference input signal and the frames of the outputs.

In some cases the timing reference is provided by another signal or clock, and the incoming signal needs to have been already synchronized to that clock. AES11 defines a specification for this sort of synchronization. It also covers synchronization of the digital audio interface with video.

Synchronization by embedded clock

The simplest form of synchronization of a single device is when there is only one input signal and that signal is used as the timing reference. This is often not thought of as synchronization because it is implicit to the operation of such equipment, such as outboard stereo DACs or stereo digital recorders (DAT, DCC, or CD-Audio).

If this input also has an output associated with it, then according to AES11 that output should be aligned so that the time difference between frame starts at input and output is less than 5% of a frame period, or 6.4 UI. This is shown in Figure 16, with the circle representing the possible phases of relative input and output frame timing. In this picture the distance around the perimeter of the circle corresponds to one frame.




Figure 16. Synchronization to the input signal.

DARS

The digital audio reference signal, or DARS, is an AES3 signal that is used for timing purposes rather than for carrying audio data. This signal can be fed from the clock master device to other devices—which would be synchronization slave devices—which need to be synchronized to each other

or to the clock master. For example, the clock master may be a digital mixer. The slaves may be the various source devices that feed the inputs of the mixer, such as tape and hard disk recorders and outboard analog-to-digital converters.

These slave devices also need to meet the AES11 output tolerance alignment specification of $\pm 5\%$ of a frame. As a result, the signals from the slave devices are appropriately aligned to the internal timing of the mixer so that there is no ambiguity about which frames are associated with the same sample time.

Input data alignment

Consider the case where an input signal is used, but it is not the synchronization source. For example, we might have a digital mixer using a DARS reference from a master synchronization clock, and several other input signals that need to be timed together. The data from the input signals need to be processed in synchronization with each other and with the DARS signal so that sample data corresponding to the same sample time are processed together.

In this situation it is assumed that the input signals are all at the same sample rate and have been synchronized by a DARS. If any are at slightly different rates then a re-synchronizing sample rate converter would be required.

The timing of the arrival of the data frames from each input signal will determine which frames are aligned together when processed. If the timing is closely matched there is no ambiguity, but if one of the input signals is slightly misaligned that produces a problem.

For this example, consider that the data from each input signal is received and decoded and briefly held in a buffer store. At a time determined by the mixer’s own clock (which is derived from the DARS) this buffer store is transferred to another store, or “read.” This defines the boundary between times when an input data word corresponds with one sample or the next.

An ambiguity in frame alignment can occur if the input signal arrives just at the time when the mixer is reading the input data buffer. If the new frame data for input sample N has been decoded and then loaded into the mixer input data buffer just before mixer sample M is read, then the input sample number and mixer sample number are the same. However, if input sample N arrives a few microseconds later, then input sample $N-1$ is used for mixer

Consumer format channel status fields								
0	Pro/con = 0 bit 0	Non-audio = 0 1	Copyright 2	Emphasis 3 4 5			Channel status mode = 00 6 7	
1	Category code bit 8 9 10 11 12 13 14 15							
2	Source number bit 16 17 18 19				Channel number 20 21 22 23			
3	Sampling frequency bit 24 25 26 27				Clock accuracy 28 29			
4	Word length bit 32 33 34 35				(Future original sampling frequency?) 36 37 38 39			
5-23	Reserved bits 40-191							

Consumer channel status

The tables to the left apply if bit 0 is set to "0"—consumer application. The definitions for bytes 1 to 23 only apply if bits 1, 6 and 7 are set to "0," for linear PCM audio and channel status mode 0. In practice the standard requires that the value of bits 6 and 7 are always set to "0" until any future revision defines them.

Note that the bit fields are shown with the earliest, or lowest numbered bit, to the left. As the format is LSB first this is the opposite way around to conventional binary notation which shows the MSB to the left.

sample *M*. This produces a time error of one sample for that input.




Figure 17. Example of input timing ambiguity with a DARS reference.

Even worse is the situation where the input sample arrives so close to the moment that the input buffer is read, that a small amount of jitter causes a fluctuation of states between a delay of one sample, and no delay at all. This is shown in Figure 17. This could result in the missing and repeating of input samples each time the data arrival phase crossed the buffer "read" phase.

The AES11 rules address this problem with the combination of input and output alignment tolerance. The output tolerance has already been mentioned. The input tolerance requires that the receiver should correctly process an input that has arrived with a timing that is within 25% of a frame period to the timing of the reference. This range is shown in Figure 18.

A receiver that needs to support the DARS synchronization mode should be designed with the input buffer read time opposite in phase to the ideal phase alignment determined by the timing of the DARS.




Figure 18. AES11 input alignment tolerance.

AES11 requires that a receiver should treat synchronized input data as being

Consumer format channel status field interpretations			
Bits	label	interpretation	
0	pro/con	0: consumer; 1: professional format	
1	non-audio	0: suitable for conversion to analog audio using linear PCM 1: not suitable	
2	copyright	0: asserted; 1: not asserted	
3-5	emphasis	000: Emphasis not indicated 100: emphasis—CD-type	
6-7	channel status mode	00: mode zero; other values reserved	
8-15	category code		
16-19	source number	(bit 16 is LSB)	
20-23	channel number	(bit 20 is LSB)	
24-27	sampling frequency	0000: 44.1 kHz 0100: 48 kHz 1100: 32 kHz	
28-29	clock accuracy	10: Level I, ±50 ppm 00: Level II, ±1000 ppm 01: Level III, variable pitch shifted	
30-31	reserved		
32	word length (field size)	0: Maximum length 20 bits 1: Maximum length 24 bits	
33-35	word length	000: not indicated	if bit 32 = 1 if bit 32 = 0
		101: 24 bits	not indicated
		001: 23 bits	20 bits
		010: 22 bits	19 bits
		011: 21 bits	18 bits
		100: 20 bits	17 bits
			16 bits
36-39	reserved		
40-191	reserved		

Professional format channel status

The table below and the table on pages 8 and 9 show the professional implementation of the channel status data. Note that the bit fields are shown with the earliest, or lowest numbered bit, to the left. As the format is LSB first this is the opposite way around to conventional binary notation which shows the MSB to the left.

Byte	Professional format channel status fields						
0	Pro/con = 1	Non-audio	Emphasis		Lock	Sample frequency	
	bit 0	1	2	3	4	5	6 7
1	Channel mode			User bit management			
	bit 8	9	10	11	12	13	14 15
2	Use of auxiliary mode sample bits		Source word length		Alignment level		
	bit 16	17	18	19	20	21	22 23
3	Channel Identification for multichannel application						
	bit 24	25	26	27	28	29	30 31
4	DARS		Sample frequency (fs)				fs scaling
	bit 32	33	34	35	36	37	38 39
5	Reserved						
	bit 40	41	42	43	44	45	46 47
6	Alphanumeric channel origin data (first character)						= 0
	bit 48	49	50	51	52	53	54 55
7	Alphanumeric channel origin data						= 0
	bit 56	57	58	59	60	61	62 63
8	Alphanumeric channel origin data						= 0
	bit 64	65	66	67	68	69	70 71
9	Alphanumeric channel origin data (last character)						= 0
	bit 72	73	74	75	76	77	78 79
10	Alphanumeric channel destination data (first character)						= 0
	bit 80	81	82	83	84	85	86 87
11	Alphanumeric channel destination data						= 0
	bit 88	89	90	91	92	93	94 95
12	Alphanumeric channel destination data						= 0
	bit 96	97	98	99	100	101	102 103
13	Alphanumeric channel destination data (last character)						= 0
	bit 104	105	106	107	108	109	110 111
14	Local sample address code (32-bit binary, LSW)						
	bit 112	113	114	115	116	117	118 119
15	Local sample address code (32-bit binary)						
	bit 120	121	122	123	124	125	126 127
16	Local sample address code (32-bit binary)						
	bit 128	129	130	131	132	133	134 135
17	Local sample address code (32-bit binary, MSW)						
	bit 136	137	138	139	140	141	142 143
18	Time of day code (32-bit binary, LSW)						
	bit 144	145	146	147	148	149	150 151
19	Time of day code (32-bit binary)						
	bit 152	153	154	155	156	157	158 159
20	Time of day code (32-bit binary)						
	bit 160	161	162	163	164	165	166 167
21	Time of day code (32-bit binary, MSW)						
	bit 168	169	170	171	172	173	174 175
22	reserved			Reliability flags			
				bytes 0-5	bytes 6-13	bytes 14-17	bytes 18-21
	bit 176	177	178	179	180	181	182 183
23	Cyclic redundancy check character (CRCC)						
	bit 184	185	186	187	188	189	190 191

sampled at the same instant, if the frame start is aligned to the DARS frame start with an error of less than 25% of a frame period. This timing offset tolerance allows for a chain of devices that are synchronized using the signal embedded clock (rather than a DARS) and therefore adding up to 5% of a frame of error for each device, and also for other timing errors.

A good receiver design can go further than this. It could use hysteresis in the region of non-compliant input timing and take away the risk of any particular timing relationship resulting in the dropping and repeating of samples. The $\pm 25\%$ rule mentioned above allows for hysteresis in the other 50% of the timing circle. This could be implemented to ensure that if the relative input alignment drifts past the critical phase, a sample of input data is not lost or repeated until the timing is up to 75% of a frame away from the nominal ideal. If that occurs and the alignment drifts in the other direction, then the correction in the other direction would not occur until the error had reduced down to 25% from the nominally ideal timing. This will then give a tolerance-to-timing wander of as much as 50% of a sample frame, even if the source has a worst-case misalignment of 180 degrees to the correct (reference) phase.




Figure 19. Provision of input buffer hysteresis to improve response to "wander."

STANDARDS

There are several published standards documents that either define basically the same digital audio interface as AES3, or the similar consumer-targeted equivalent, in IEC60958-3. There are also standards that are used in conjunction with these interfaces.

IEC60958:1989 (previously known as IEC958:1989)

This has been replaced by the multi-part document, IEC60958-n. It defined both the professional and consumer applications and the two-connector types for electrical connection. By accident it did not require that the professional format used the XLR and the consumer format used the coaxial connection.

Professional format channel status field interpretations

Bits	label	interpretation																					
0	pro/con	0: consumer; 1: professional format																					
1	non-audio (or, more accurately, "not linear PCM")	0: audio data is linear PCM samples 1: other than linear PCM samples																					
2–4	emphasis	000: Emphasis not indicated 100: No emphasis 110: CD-type emphasis 111: J-17 emphasis																					
5	lock	0: not indicated 1: unlocked																					
6–7	sampling frequency	00: not indicated (or see byte 4) 10: 48 kHz 01: 44.1 kHz 11: 32 kHz																					
8–11	Channel mode (SCDSR = single channel double sample rate)	0000: not indicated (default to 2 ch) 0001: 2 channel 0010: 1 channel (monophonic) 0011: primary / secondary 0100: stereo 0101: reserved for user applications 0110: reserved for user applications 0111: SCDSR (see byte 3 for ID) 1000: SCDSR (stereo left) 1001: SCDSR (stereo right) 1111: Multichannel (see byte 3 for ID)																					
12–15	user bit management	0000: no indication 0001: 192-bit block as channel status 0010: As defined in AES18 0011: user-defined 0100: As in IEC60958-3 (consumer)																					
16–18	use of aux sample word	0000: not defined, audio max 20 bits 0001: used for main audio, max 24 bits 0010: used for <i>coord</i> , audio max 20 bits 0011: user-defined																					
19–21	source word length	<table style="width: 100%; border-collapse: collapse;"> <tr> <td style="width: 10%;"></td> <td style="width: 45%; text-align: center;">if max = 20 bits</td> <td style="width: 45%; text-align: center;">if max = 24 bits</td> </tr> <tr> <td>000:</td> <td style="text-align: center;">not indicated</td> <td style="text-align: center;">not indicated</td> </tr> <tr> <td>001:</td> <td style="text-align: center;">23 bits</td> <td style="text-align: center;">19 bits</td> </tr> <tr> <td>010:</td> <td style="text-align: center;">22 bits</td> <td style="text-align: center;">18 bits</td> </tr> <tr> <td>011:</td> <td style="text-align: center;">21 bits</td> <td style="text-align: center;">17 bits</td> </tr> <tr> <td>100:</td> <td style="text-align: center;">20 bits</td> <td style="text-align: center;">16 bits</td> </tr> <tr> <td>101:</td> <td style="text-align: center;">24 bits</td> <td style="text-align: center;">20 bits</td> </tr> </table>		if max = 20 bits	if max = 24 bits	000:	not indicated	not indicated	001:	23 bits	19 bits	010:	22 bits	18 bits	011:	21 bits	17 bits	100:	20 bits	16 bits	101:	24 bits	20 bits
	if max = 20 bits	if max = 24 bits																					
000:	not indicated	not indicated																					
001:	23 bits	19 bits																					
010:	22 bits	18 bits																					
011:	21 bits	17 bits																					
100:	20 bits	16 bits																					
101:	24 bits	20 bits																					

IEC60958-1, -3, and -4

The revision of IEC958:1989 involved splitting the standard into three parts. Part 1 covers the aspects common to both consumer (which is in

part 3) and professional (part 4) applications. As the document has a different reference it is also

Edition 1—which may be confusing.

AES3-1992

The primary definition of the professional format is in this document. This undergoes regular revision by amendment or new edition. It is possible for interested parties to contribute to this process by joining the working group on digital input/output interfacing, SC-02-02. Further information on joining AES standards working groups is available at <http://www.aessc.aes.org>.

EBU 3250 (Ed. 2, 1992)

This document has been produced by the European Broadcasting Union (EBU). It is similar to AES3-1992 (without the amendments) apart from one key difference—the EBU document specifies that transformers shall be used between the cable connection and the receiver and transmitter electronics. Transformers are optional for AES3.

ITU-R BS647-2 (1992)

This is very similar to EBU3250. The International Telecommunications Union (ITU) is an intergovernmental organization that is part of the UN.

IEC60958-4 (Ed. 1, 1999)

This part of IEC60958 defines the professional interface. At the time of writing (early 2001) it is similar to AES3-1992 with amendments 1 (1997) and 2 (1998) but not amendments 3 or 4 (both 1999); the key difference is that it does not support sampling frequencies above 48 kHz. There is an amendment in process within IEC to rectify this.

Technical Report IEC60958-2:1994 (or IEC958-2:1994)

This document is not a standard. The specification describes a method of carrying software information in the channel status stream of the consumer application of IEC60958:1989. It uses the setting of the channel status mode field in byte 0 of the channel status block to distinguish this use of the channel status block. Originally it was proposed as an amendment to the IEC958:1989 standard but was rejected. With the conversion of that standard to a three-part standard, with parts called IEC60958-1, IEC60958-3, and IEC60958-4, this document appears—at first sight—to be part 2 of a 4-part standard. That is not the case.

IEC60958-3 (Ed. 1, 1999)

This part of IEC60958 defines the consumer interface in all respects except the optical interface. At the time of writing (early 2001) it does not support sampling frequencies above 48 kHz. There is an amendment in process within IEC to rectify this.

Professional format channel status field interpretations (Cont.)

Bits	label	interpretation
22–23	alignment level	00: not indicated 01: –20 dB FS 10: –18.06 dB FS
24–31	channel identification	if bit 31 = 0 then channel number is 1 plus the numeric value of bits 24–30. if bit 31 = 1 then bits 4–6 define a multichannel mode and bits 0–3 give the channel number within that mode.
32–33	digital audio reference signal (DARS)	00: not a DARS 10: DARS grade 2 (+ / –10 ppm) 01: DARS grade 1 (+ / –1 ppm)
35–38	sampling frequency	0000: not indicated 1000: 24 kHz 0100: 96 kHz 1001: 22.05 kHz 0101: 88.2 kHz 1101: 176.4 kHz 1111: User defined
39	sampling frequency scaling	0: no scaling 1: apply factor of 1 / 1.001 to value
48–79	alphanumeric channel origin	four-character label using 7-bit ASCII with no parity. Bits 55, 63, 71, 79 = 0.
80–111	alphanumeric channel destination	four-character label using 7-bit ASCII with no parity. Bits 87, 95, 103, 111 = 0.
112–143	local sample address code	32-bit binary number representing the sample count of the first sample of the channel status block.
144–175	time of day code	32-bit binary number representing time of source encoding in samples since midnight
176–183	reliability flags	0: data in byte range is reliable 1: data in byte range is unreliable
184–191	CRCC	00000000: not implemented X: error check code for bits 0–183

AES-3id 1995 and SMPTE 276M-1995

These two documents both define a variant of AES3 that is transmitted over 75 Ω coaxial cable at a level of 1 V (peak to peak). The impedance and level are chosen to be compatible with broadcast analog video interfacing and allow the use of some of the same cabling and routing infrastructure.

The two specifications are different. The SMPTE specification has tighter tolerances for some parameters and intended for use with dedicated

interfaces on equipment for high performance. AES3-id has more relaxed specifications that permit use with passive converters between the 110 Ω balanced and the 75 Ω coaxial formats.

AES11-1991 (Synchronization)

This standard defines rules to be followed to ensure synchronization of digital audio equipment together and with video. A special AES3 signal is used to distribute a timing reference from the synchronization (clock) master

to all the other (slave) devices in the synchronized system. This timing reference is called a Digital Audio Reference Signal (DARS). **Ap**



Julian Dunn took degrees in Astronomy and then Medical Electronics at London University, where he first became interested in signal processing. After graduating in 1984 he joined the BBC Designs Department, where he designed digital audio equipment as part of work for BBC Radio in prototyping equipment for use with the new digital audio recorders, mixing consoles and transmission systems.

After a year working at the Mullard Radio Astronomy Observatory in Cambridge, Julian joined Prism Sound as a consultant, once again designing digital audio equipment. In 1998 he formed his own digital audio design company, Nanophon. Nanophon provides specialist consultancy in digital audio conversion, DSP software and hardware, digital audio interfacing and clock recovery systems.

Julian has presented technical papers to AES conferences and conventions on various topics in digital audio, and has been contributing to the work of the AES standards digital audio subcommittee since 1991. You can find Julian at <http://www.nanophon.com>.

APWIN Version 2.14 Now Available

APWIN Version 2.14 replaces version 2.11. This is primarily a maintenance release that incorporates several bug fixes and improved drivers for the PCI and PCMCIA interface cards.

Version 2.14 will run correctly on the Japanese version of Windows. Improvements in compatibility between interface cards and various operating systems have been made; in particular, the PCMCIA interface (with new firmware) now works with Windows 2000. These changes are detailed in the Compatibility Chart on our web site.

There will not be a general distribution of Version 2.14, but you can download it from our web site at <http://audioprecision.com/software/updates/index.htm>. **Ap**

New Sales Support Engineer

David Post



Meet David Post, Audio Precision's new Sales Support Engineer. Dave will be working under Director of Sales and Marketing Tom Mintner to develop new demos and otherwise help with trade shows, partner with our sales reps in visiting customers, explore new market opportunities, and add support to sales administration.

Dave started his career in the Navy, working in avionics and metrology. He then moved to General Dynamics as a quality control engineer in a metrology lab, followed by sales engineer positions for companies such as Fluke, LeCroy, and Nicolet, and finally spent nine years building a manufacturer's rep business in Fort Worth.

Dave and his wife are both native Texans, and have an adopted son who is now four years old. When not on the daddy track, Dave plays electric bass, enjoys his home hobby recording studio, and avidly reads classical books. At least he did, before beginning to discover the wonders of the Pacific Northwest. **AP**



New TECHNOTES Available

.WAV File Import Utility

Measurement Microphone dBspl Calibration

The sound pressure level of loudspeakers and other acoustic transducers is measured in absolute terms. The standardized unit is dBspl, or dB sound pressure level. TECHNNOTE 28 describes how to set up and use APWIN's dBr unit as dBspl. This procedure applies equally to the analog analyzers of System One, System Two, and System Two Cascade.

To use APWIN's dBr unit as dBspl, the 0 dBspl value for the specific microphone must be computed and entered into the REF dBr field near the bottom of the Analog Analyzer panel. Two methods are described: you can either compute the 0 dB value, or use a microphone acoustic calibrator and adjust the system for the expected reading.

The details of the computation, plus an APWIN procedure to automate them, are available on our Web site: audioprecision.com/publications/TECHNOTES/index.htm and download TECHNNOTE 28. **AP**

This APWIN utility allows you to load any .wav files compatible with Microsoft RIFF format into System One, System Two, or System Two Cascade, and analyze them in the frequency or time domain. The utility takes into account the selected post-processing digital analyzer (DSP Analyzer) program and its panel settings and creates the required waveforms to be loaded into the Acquisition and Digital Generator buffers if required.

The sample rate and resolution (number of bits) are automatically determined from the source .wav file and the Digital Input/Output panel settings are updated accordingly. The import process requires that the .wav file be converted to a file compatible with Audio Precision DSP hardware. This process requires information about the Digital Analyzer program to create the file correctly.

The sample rate of the .wav file must be supported by the selected post-processing digital analyzer (DSP Analyzer) program or an error message will inform the user that the sample rate is not supported. The utility will not do sample rate conversion.

The acquire length is used to define the file size of the waveform loaded into the acquisition buffer and if needed the

size of the file loaded into the generator buffer. If the .wav file length is shorter than the desired acquisition buffer length, the waveform is duplicated until the desired length is obtained. If the waveform is longer than the desired acquisition buffer length, the transferred portion of the waveform starts at the beginning of the .wav file and ends at the acquisition buffer length.

A DSP post processing digital analyzer program is required to analyze imported .wav files. For each of the systems, we have real time and post processing DSP programs that provide different measurement capabilities. A real time program constantly analyzes data as it is supplied to the DSP and calculates and displays measurements on the panels in much the same way that readings are returned from the analog sections of the hardware. In a post-processing program, the incoming signal is digitized for the defined acquisition length then processed to produce an FFT from which measurements are returned to APWIN for display.

This procedure is available at our Web site. Just go to audioprecision.com/publications/TECHNOTES/index.htm and download TECHNNOTE 30. **AP**

Date	Event	Location
April 23-26, 2001	NAB	Las Vegas
May 12-15, 2001	AES Europe	Amsterdam, Holland
September 21-24, 2001	AES New York	New York City

Power Users:

How to turn off Autofill

Power users, are you tired of those pesky AutoFill menus that appear when typing Ap.xxx.yyy.zzz in the procedure Editor? Disable this feature by unregistering Apbauto32.dll, using the following procedure.

- Through Windows Explorer, confirm the location of the file Apbauto32.dll. Typically, it's located in C:\Program Files\Audio Precision\Apwin200, and our example assumes this is the case.
- Open a DOS window, then enter the following, including the quotes and spaces:
cd "C:\Program Files\Audio Precision\Apwin200"
- Press **Enter**. Your prompt should change to reflect this path.
- Next, type:
C:\Windows\System\Regsvr32 /u Apbauto32.dll

- Press **Enter**. You should see the following message window:



- Click **OK**, then type **exit** to close the DOS window.

To re-register Apbauto32.dll, follow the same procedure but without the "/u" switch. [Ap](#)

Much More about Monitor Outputs

Monitoring Analog Signals

The users manuals for the System One, System Two, and System Two Cascade briefly discuss the monitor outputs, but have left some questions unanswered. This article addresses the details of the Analog Monitor Outputs.

The analog monitor connections provide buffered outputs for an oscilloscope or amplifier/speaker without needing to connect them to the signal being measured. This prevents interference with an accurate audio measurement that would occur if a parallel connection to an oscilloscope or power amplifier were used. The monitor outputs are not calibrated to be used for accurate measurements.

The analog monitor sources are from either the Level meters or the Reading (Function) meters after internal ranging for optimum measurement level. The result is that they typically have a maximum amplitude of about 3.5 Vp-p before a range change brings them down to a lower level—about a 2-to-1

monitor amplitude change. At very low levels of input signal, the lowest range will be in use so the monitor signal output may also be lower in amplitude.

If AutoRanging is not checked, the monitor amplitude will be a specific ratio to the analyzer input level and will depend on the fixed range in use. If the fixed Level meter range is 1.2 V, the monitor outputs have a 1:1 monitor voltage to input voltage ratio. The monitor output is halved for each doubling of input range above 1.2 V: the 2.5 V range ratio is 1:2, 10 V range 1:8, etc. Below the 1.2 V range, the monitor output doubles for each halving of the input range: the 600 mV range ratio is 2:1, the 160 mV range ratio is 8:1. The most sensitive Level Meter range on System Two and Cascade is 40 mV. At this range, the monitor output to input voltage ratio is 32:1. The most sensitive Level meter range on System One is 80 mV.

When AutoRange is not checked in the Level meters, the indicated value in the Range field is the maximum peak amplitude that is within the specified accuracy of the Level meter (meaning the reading is correct). In the Function meter (or Reading meter—we use these terms interchangeably), when AutoRange is unchecked, the amplification into the Function meter is shown. Depending on the active Function, this may be a gain of 1 (or 0 dB) up to a gain of 1024. This provides better accuracy, for example, with a distortion measurement after the notch filter has removed the fundamental.

The source impedance of the analog monitor outputs is 600 Ω. This is adequate for monitoring devices and protects the monitor if the output is shorted. Remember, these are not calibrated signal source but are only for signal monitoring use. [Ap](#)

Mike Wolfe

Applications Support Engineer

Perhaps you've already met Mike Wolfe. That's because he's a member of Audio Precision's Technical Support Group, where he fields customer inquiries and teaches training classes. Mike works in audio because he likes music. He comes by it naturally; his father was a composer who attended Juilliard School of Music on the GI Bill. Mike likes his music with the electronics and recording media as unobtrusive as possible.

Mike says he's been around "since electrons came in bottles." He started his career as a research assistant in New Mexico for MIT, setting up arrays of scintillation detectors to study cosmic rays. He's been in audio since 1970, as



Mike and his two grandsons, Jackson (L) and William (R) at William's first birthday party.

a recording engineer, designing and maintaining professional audio studios, or working in FM stations or in the audio end of TV broadcasting. Small wonder he should now be working with the equipment that defines Audio Precision.


But Mike is a person of many other interests; "200 years worth," he says. His wife, Jill, is a semiprofessional photographer, and they like to travel the back roads of the Western U.S. Their junkets usually include about 80 lbs of cameras and equipment, ranging from a 6 x 7 cm Pentax, to an old Rollei that he loves, to several 35mm Nikons (Mike scorns automatic cameras). Mike is an avid birdwatcher, and captures birds on film. They are also likely to return from the outback with many pounds of rocks, as his wife is also an amateur geologist.

Mike and Jill have three daughters, two of whom are married and have presented them with a grandson each.

Mike names Alexander Calder, the sculptor, and Richard Feynman, the theoretical physicist, as two of his heroes. Given the opportunity, Mike says, he would like to devote more time to sculpting, utilizing what he learned in a welding sculpture class



This view of the Cape Mears Lighthouse was taken on a recent trip to the Oregon coast. Completed in 1890, the light could be seen 21 miles out to sea. It was retired by a new lighthouse in 1963.




Mike often finds unique views of impressive structures.



Bird watching and photography are natural companions for Mike.

recently. Or enlarging their collection of art of the Southwest and Plains Native Americans. Or perhaps Mike would take up a new vocation entirely: a new employee recently said of Mike, "He looks like Santa Claus." **Ap**

Audio
precision®



Audio Precision
P.O. Box 2209
Beaverton, Oregon 97075-2209
Tel: 503-627-0832 Fax: 503-641-8906
US Toll Free: 1-800-231-7350
Email: techsupport@audioprecision.com
Web Site: audioprecision.com