AES3 Channel Status Revisited

MARK YONGE MIBS continues and expands the discussion on the AES3 Interface.

I was very happy to see Mike Law's article, *AES3 Status Information – What is it For?* in the April/May issue of *Line Up*. Any 20 year-old high-tech standard might reasonably be questioned, and any considered critique of this little-understood aspect of the world's best-known professional digital audio interface is useful (and sufficiently rare to be noteworthy in itself). I'm writing here, as a member of the IBS, because I felt I ought to be able to put some of Mike's points into perspective. I did a little research and I'm also grateful to Robin Caine whose knowledge of the subject, and the history, is close to encyclopaedic. It turns out that some of the history is relevant.

The AES3 standard for a professional two-channel digital audio interface emerged from the general discussions about digital audio in the late 1970s and early 1980s. There were a number of interested parties. Some were interested in interfaces for CD players and future consumer equipment; others, like the EBU, were interested in professional applications in broadcast; a third group were interested in professional equipment for music recording. The Audio Engineering Society provided a forum for these discussions.

Interface Divergence

Many interface issues were sufficiently general that solutions could be adopted by all parties but, when it came to the use of the Channel Status data block, the professional and consumer interests diverged. The EBU and the AES, concentrating on professional applications felt strongly that the channel status data should be used to support a range of options potentially important to professional users. The consumer group had a different agenda.

As a result, not one but two standards emerged. The frame and subframe of the data stream, the block structure for additional data, and the PCM-coded audio data itself, was similar in both cases. The AES and the EBU defined a set of Channel Status parameters with checksum protection to avoid corrupted status data. IEC 60958-3 (aka S/PDIF) defined a simpler scheme, without checksum protection, to support copyright management. Let's look again at the important parts of that block of 192 bits of status data that are repeated every 192 samples.

Byte 0, bit 0: Consumer/Professional mode

How would you tell these two similar-but-different formats apart? It was agreed that the first bit in the first byte of the Status Data would indicate whether the content of that Status Data conformed to the professional (AES3) or consumer (IEC 60985-3) standard. This is arguably the most important bit in the whole 192-bit Channel Status block. Hopefully, we've now got to: Confusion 0%, Importance 100%, Usefulness 100%

Byte 0, bit 1: Linear or Non-linear samples

It's linear PCM audio unless you set this bit. This has never been controversial.

Byte 0, bits 2 to 4: Pre-emphasis

Twenty years on, the option for including HF emphasis in PCM-coded audio does seem a little silly. However, at the beginning, many practical ADCs and DACs were still 13-bit and 14-bit devices and the noise floor was very apparent. HF emphasis was a technique that had been used in FM radio, for example, to 'borrow' HF headroom to improve HF noise performance. However, this reckoned without new forms of music with lots of HF energy – there was no headroom left to borrow – so emphasis didn't pay off! Now, with competent converters of 20 bits and higher, the technique is only a historical curiosity.

Byte 0, bit 5: Nothing/Unlocked

This bit does not and cannot mean Locked vs Unlocked. The fact that the transmitting device feels cosy about being locked has little meaning for the receiving device that may not have access to the same clocks. However, setting this bit could be useful in planned situations to indicate that a device had lost its reference, including the Digital Audio Reference Signal (DARS) unit which might become unlocked from its parent video reference.

Byte 0, bits 6 to7: Sampling Frequency

The need to state a sampling frequency was more necessary before dedicated silicon came into use and everyone had to design their own logic (remember TTL?). With the use of over-sampling DACs there was easier indication from the clock rate than from channel status. It was useful when a digital reel-to-reel was started at the wrong speed, giving 48kHz when the original was 44.1kHz. Then the channel status would set the error lights flashing.

Byte 1, bits 0 to 3: Channel Mode
This illustrates nicely why there are two separate channel status streams, one in each channel. AES3 is a two-channel interface and there is no guarantee that both channels will be similar, or even related. For example, one could be PCM-coded audio, the other some different type of data. In router management the two channels can be, and often are, split and routed independently.

Byte 1, bits 4 to 7: User Bit Management
User data has not been used to any great extent to date, however the apparent thirst for metadata in the broadcasting community at least, suggests that its use will increase. For example, one project currently being developed (AES-X111) proposes that User Bits may carry a unique identifier for the associated audio stream, potentially handy for relating any amount of external metadata to a specific audio source. Press the Red Button now!

Byte 2, bits 0 to 2: Aux Bits
The BBC formalised a use for Aux bits, though they possibly never used it. There was also a suggestion of putting the four protection bits from a Hamming 11,7 code to error-correct the seven MSBs of the sample into the Aux bits – errors below the first seven MSBs have lower perceptibility than concealment mechanisms. This scheme was never formalised.

Byte 2, bits 3 to 5: Word Length
The design of the scheme meant that failure to read these status bits was never catastrophic; even the accidental playing out of Aux data would be at a very low level. In fact the default of 20 bits was safe, if not purist.

Byte 2, bits 6 to7: Alignment Level
The IABM asked AES Standards to implement these two conditions because of potential problems with videotape interchanges between European and US users. The AES standardised the channel status flags simply to indicate the use of two established practices standardised elsewhere (EBU R68 and SMPTE RP155).

Byte 23, bits 0 to 7: CRCC
This was never meant to indicate transmission errors, still less that there was a problem with the audio. It merely shows that a block of channel status data has been corrupted; by editing or switching for example.

Implementation

One of the biggest issues in completing the 1985 standard was the need to specify three levels of implementation. The 'Minimum Level' simply read Byte 0 bit 0 to determine whether the bitstream was professional or consumer. This requirement was demanded by manufacturers who were unwilling to incur the cost of implementing real channel status because they were actually thinking in consumer terms. The committee was forced to include this 'Minimum level' of channel status so that professional receivers would accept this audio. 'Standard level' was what most users required, with 'Enhanced' implementation available for those with such needs. The items mentioned above comprise the AES3 'Standard Implementation.' All other bytes are optional, and basic functionality should be independent of them.

The Future

In the early days, the structure of channel status was largely formed by the difficulty of designing logic to perform the decoding. Many early decoders 'cherry-picked' the channel status bits of interest and ignored the rest – rather than parsing the bitstream properly starting at byte 0 – causing some obvious problems.

The AES is not a policeman – AES standards are voluntary – and requirements in the standard will never stop silicon manufacturers selling whatever they can. However, if a chip or a piece of kit claims 'AES3' but doesn't comply with the standard, someone ought to point this out – to the supplier at least.

AES3 may support more complexity than any one application may need, and certainly the predictions of 1985 have led to some redundant features and a wish-list of alternative definitions. However, much of it is used and the standard is still being expanded to support new requirements such as higher audio sampling frequencies and modern metadata schemes. Not bad for a 20-year-old digital document, and it'll stay useful for as long as good people help to keep it up to date and users appreciate its usefulness and dependability.

jbs

Further Reading

In addition to *AES3-2003: Serial transmission format for two channel linearly represented digital audio data*, the AES publishes two information documents intended to help those implementing the AES3 interface. They are, *AES-2id: Guidelines for the use of the AES3 interface*, and *AES-3id: Transmission of AES3 formatted data by unbalanced coaxial cable*. All these documents are available for download at www.aes.org/publications/standards/

The AES Historical Committee has put together an archive of early digital standardisation reports which will illuminate the early development process. Go to www.aes.org/aeshc/ and click on *Digital Audio Engineering Standardization History at the AES*.

AES Standards general information is at: www.aes.org/standards/ The *IEC 60958 (part 3) Digital audio interface: Consumer applications* is available from www.iec.ch/