Broadcast Audio
Custom Manufacture
Design and Consultancy

Acrone Ltd
5 Bristol Way
Stoke gardens
Slough, SL1 3QE

Tel. 01753 579524
Fax. 01753 577981

Comparison between 24bit PCM and 20+4bit scaled PCM

Professional Audio systems are based upon a 24bit signal, and the existing hardware and chip sets are based upon 24bits.

The data paths present in AES3, MADI and AOIP are currently based upon a 24bit architecture. AES3, MADI and many data paths inside equipment are unable to cope with signals at a wider bit-width, so it is useful to see if an alternative to 24bit PCM is possible.

These experiments were undertaken in December 2014 by BCD Audio to explore how a 20bit PCM + 4 bit signed scale factor would compare in performance.

TRANSMITTER

A transmitter starts with a 24bit linear PCM signal, for example from an AtoD converter.

Currently any digital gain required at the transmitter will create values in excess of 0dBFS, 24bit, so will clip in a 24bit architecture.

Also if a digital loss is required at the transmitter, the signal will be truncated to a lower resolution before transmission.

Many systems will be using a DSP or FPGA capable of more than 24bit working, so that the bottleneck is now the 24bit transmission path only.

The system proposed here adds one simple software or hardware module to SCALE the signal before transmission.

This module computes the number of sign bits present in the signal, and provides a scale factor as a result.

A 4 bit scale factor is a good match to the dynamic range of Audio signals, and allows a maximum shift of 7 bits up and 8 bits down.

A scale factor of 0 results in no shift, and is equivalent to 0dBFS with 24bit linear PCM.

The signal is shifted by this number, and then truncated to suit the output requirement.

A suitable truncation width is 20bit, and shift is 4bit, which still has a total of 24bits, but now has a dynamic range of 186dB.

Digital overload at 0dBFS is no longer a problem, as the signal can reach +42dBFS before overload.
Low level truncation is no longer a problem, as the signal can now resolve to -168dB.

For the best performance, a transmitter will calculate the scale factor on every sample, with no latency.

Simpler transmitters could use a fixed scale factor, or a smaller number of scale factors.

RECEIVER

A 20bit PCM + 4bit scale factor signal uses a signed 4bit binary scale factor in addition to the the PCM signal.
A receiver can pick up the 24bits, and separate out the two fields.

There are no latency or ambiguity issues, as the scale factor always belongs to the same PCM sample.

The receiver uses the scale factor as a simple shift factor, and scales the signal up and down depending on the factor.

Many processors contain a ' barrel shifter' that can execute this operation quickly.
The operation is also simple enough to implement in an FPGA.

The scaled signal restores the original PCM signal present at the transmitter, and will create a linear PCM signal at 35bit resolution.

DISTORTION

The contentious issue is whether a 20bit PCM + 4bit scale factor signal is suitable for a new system or not.

Distortion measurements of Linear PCM at spot levels and differing bit widths were made.

Tests were conducted on a BCD Audio 24bit DSP card, in simple pass through on the left signal.

The bit width on right signal was changed by a small amount of DSP code, to simulate truncation and scaling effects.

The test generator and receiver used was a Prism Dcope 3.

Results are shown for a 497HZ sine wave signal, and some spot levels.

The experiment was repeated for other frequencies and levels and the results are the same.

SIMPLE TRUNCATION

To get an idea how simple truncation affects a signal, the following was measured.

A rounding algorithm was applied to the 20bit version , and no appreciable difference was found.

This was surprising, as I expected to see some difference, and might be attributable to the maths routines in the Dscope.

Signal width	0dBFS	-6.02dBFS	-12.04dBFS	-18.06dBFS	-24.08dBFS
16bit	0.00117%	0.0023%	0.0047%	not measured	not measured
17bit	0.0006%	0.0012%	0.0023%	not measured	not measured
18bit	0.0003%	0.0006%	0.0018%	not measured	not measured
19bit	0.00015%	0.0003%	0.0006%	not measured	not measured
20bit	0.00007%	0.00015%	0.0003%	0.0006%	0.0012%
20bit + rnd	0.00007%	0.00015%	0.0003%	0.0006%	0.0012%
24bit	0.00001%	0.00003%	0.00006%	0.00012%	0.0023%

All linear PCM systems will show increasing distortion at lower levels, as the number of significant bits reduces with level.

At the lowest signal levels, the distortion will be masked by the system noise and dither used.

SCALING TO 20BIT

A simple routine was added , to simulate the effect of a scaled 20bit signal on a link.

The signal upper bits were checked, and if all were the same, the signal masked to the correct scale factor.

For example if the upper bits were either 11111 or 00000 no bits were masked, as the signal is already 20bit linear.

If the upper bits were 1111 or 0000 one LSB bit was removed by mask of FFFFFE

If the upper bits were 111 or 000 two LSB bits were removed.

If the upper bits were 11 or 00 three LSB bits were removed.

otherwise the bottom four LSB bits were removed by the mask of FFFFF0.

This algorithm simulates what would happen when a transmitter sends 20bit PCM with scale factors 0,-1,-2,-3, and a reciever scales the signal back to linear PCM.

In the spring 2015 real hardware was created, and the same results were measured.

Signal width	0dBFS	-6.02dBFS	-12.04dBFS	-18.06dBFS	-24.08dBFS	-48.16dBFS
20bit scale	0.00008%	0.00008%	0.00009%	0.00013%	0.00023%	0.0035%
24bit	0.00001%	0.00003%	0.00006%	0.00012%	0.00023%	0.0035%

Conclusions can be drawn here.

1. A 20bit scaled signal has worse distortion at high signal levels, but the distortion remains constant in the scaling range.

   But lets keep real! The distortion is still below 0.00008%, and a good analogue figure is better than 0.002%.

   Also the signal has a 120dB dynamic range, so only signals -110dB down will be masked with high level signals.

2. When the signal level is below -18.06dBFS, both signals are identical, so the distortion factors become identical.

   So 20bit scaled is just as good as 24bit at normal levels. No masking or truncation is present at all in the proposed signal.

3. If the source material had more than 24bits available, the distortion of the 20bit scaled version would be better than the 24bit version, as the 24bit version would truncate more than the scaled 20bit version.

   In this instance, the 20bit scaled will tend to stay at 0.00008%, where the 24bit will be worse.

4. If the source material was louder than 0dBFS, the 24bit version would clip, but the 20bit scaled version would still maintain 0.00008%

   If we allow 4bits scaling, the 20bit scaled version can scale up by +42.14dB, and scale down by 24.08dB.

NOISE

The noise floor of a 24bit PCM signal is defined by the system noise present for normal signals, and is not an issue, and is well understood.

As the scale factor is sent with every sample, the 20bit signal will be scaled back automatically, and can potentially out perform 24bit PCM.

All the receiver is doing is reversing the shift factor applied at the transmitter, so there are no latency or DC offset issues.

METADATA

The existing AES standard allows for 20 and 16bit metadata forms, so could be used here. The 20bit scaled signal would be degraded to 0.00013% distortion and would become the same as linear PCM at or below -24.08dBFS.

COMPATIBILITY WITH 24BIT PCM.

An existing 24bit PCM signal would be compatible if it was truncated to 20bit before transmission.

A 20bit + 4bit scale receiver would be compatible with 24bit PCM if a control signal was present to indicate the signal mode.

AOIP and MADI could easily provide this mode bit.

AUTODETECTION

The number of sign bits in a linear PCM signal is measurable, and is actually the loudness of the signal in multiples of 6.02dB.

As a signal gets louder, the number of sign bits reduces, and between -6.02 and 0dBFS there is only one sign bit present.

A scaled signal is normalised, so that there is only one sign bit present, apart from at really low signal levels; this is called normalisation.

Generally a 24bit PCM signal is not normalised, apart from with very loud or clipped signals.

A scaled signal is always normalised, apart from with very quiet signals, and a special code can be inserted to cover this case.

Therefore a reasonably simple algorithm can auto detect the two signals.

SUMMARY

20bit PCM + 4bit scale has 186dB dynamic range, distortion of 0.00008%, and can accept signals to +42dBFS.

24bit PCM has 144dB dynamic range, clips on overloaded signals and can have worse distortion at lower signal levels.

20bit PCM + 4bit scale could be referred to as 24S4 so that it is not regarded as an 'inferior' system!

Mike Law BCD Audio December 2014
minor updates June 2015