Trevor de Clercq
October 27th, 1998

Digital Audio Processing
E85.2600, Prof. K. Peacock

Final Report:
Consumer Benefits of the 24-bit Resolution Revolution

The professional audio industry is currently undergoing a landmark revolution. Technology that has been dominant for almost two decades is being replaced and becoming obsolete. The digital audio standard of 16-bit resolution is shifting towards the 24-bit format. Considering the usually quick turnover of technology in today's society, it is unusual that such a standard existed for so long. The endurance of 16-bit encoding, however, only attests to the importance of the current change in bit resolution. Although the only people who currently come into direct contact with 24-bit formats are recording engineers themselves, this new technology will have and already has had a significant impact on the sonic quality appreciable by the average consumer. Special encoding processes such as Super Bit Mapping, High Definition Compatible Digital, and the UV22 System allow the fidelity of high resolution 24-bit sources to be captured and recreated on 16-bit formats such as the ever present Compact Disc. Also, the appearance of 24-bit digital recorders is making 16-bit recording devices veritable dinosaurs and significantly lowering the price of 16-bit machines which were once the pinnacle of digital recording quality. The end result of this 24-bit revolution, therefore, is an easier access and greater availability of high fidelity music.

To find examples of the rapid spread of this 24-bit technology into the recording industry, one only has to open the pages of any recent trade magazine. Pro-Tools now has 24-bit capability; Lexicon effects units are operating at 24-bit levels; even editing workstations such as Sonic Solutions are offering up high resolution possibilities. At the 100th Audio Engineering Society's Show in Denmark (2 years ago), Sony was already displaying a full 24-bit production line, including the OXF-R3 24-bit digital mixing console, the 24-bit capable 48-track DASH recorder called PCM 3348-HR, and a 24-bit portable DAT with timecode. The professional audio industry is truly championing 24-bit recording as an answer to its complaints with 16-bit digital audio.

Probably one of the most important developments in the 24-bit revolution was Tascam's introduction (only a few months ago) of a 24-bit professional Digital Audio Tape recorder, the DA-45HR. The introduction of a 24-bit DAT recorder may not at first seem like a definitive change in the music industry, but the implications of this new product portended a major shift in audio technology. Without a 24-bit mixdown device, recording engineers were necessarily hesitant to accept 24-bit products. Why bother recording and mixing at a high resolution level if the products of your effort could not be faithfully preserved. With a 24-bit mixdown device, audio does not need to undergo redithering or any noise shaping process. Masters can be recorded solely over digital pathways. It is very important to note that a 20-bit DAT recorder was never released by any DAT manufacturer. In order to keep high resolution audio in the digital realm, therefore, engineers have no choice but to now embrace the 24-bit format. As much as the existence of 24-bit recorders and mixing decks influenced the development of a 24-bit DAT recorder, in turn the 24-bit DAT recorder will inspire a corollary development of 24-bit mixing boards, multi-track recorders, and computer software now that such high resolution mastering capabilities are available.

Semi-professional digital audio equipment, unfortunately, has yet to develop 24-bit technology. As of yet, the ADAT format has yet to release a 24-bit version. However, Studer has plans for a V-Eight ADAT Type II format machine which would employ ADCs from its digital 24-track DASH recorder, D827. The

DTRS format, commonly seen as Tascam's DA-88, is also still only in the 20-bit stage of recording capability. The third digital tape format, DASH (a professional format) already has numerous 24-bit models, including the previously mentioned Studer D827 and the Sony PCM-3348HR. Although AMS Neve has many 24-bit digital mixing desks, such as the Logic 2, Logic 3, and the Capricorn, this technology has not reached semi-pro mixing formats. It seems only a matter of time, however, before 24-bit technology will filter down to all levels of the audio industry. Perhaps one day even 24-bit CDs (or something comparable) will be introduced to allow consumers the full advantage of 24-bit fidelity.

But why bother with 24-bit technology anyway? Isn't 16-bit, "CD quality" already the height of sonic fidelity? With consumer formats, yes, 16-bit is the height of sound reproduction. CD quality far exceeds that of the traditional cassette tape or vinyl record in signal-to-noise ratio, distortion, lifespan, and many other areas. In the professional audio community, however, preferance is still given to the analog medium over 16-bit digital when it comes to recording music master tapes. One only has to do minimal research into the field of modern rock recording to see that analog machines are still being used on the most expensive and most successful work despite our "digital age". Is this preference the result of the analog audio industry being a more mature industry than the digital audio industry (and thus more familiar to engineers), or is analog inherently better in some ways than a purely transparent digital medium?

Engineer and AES Presenter D. Stripp explains one benefit of analog recording over 16-bit digital when he writes, "....high-frequency crushing [by analog tape] provides a subtle, mellow modification of the signal at just those high levels when the ear is getting ready to flinch at the onset of severe harmonic distortion or even clipping in other links in the chain". Furthermore, he goes on to state that, "the lack of high-frequency crushing on the digital route may cause critics to accuse it of 'hard' or 'metallic' quality". To combat this problem, the BBC used limiters and digital delays to forecast peaks in the signal and effectively prevent the these peaks from creating the "high levels when the ear is getting ready to flinch." The question for many audiophiles becomes: is there an aesthetic difference between the two modes of high-frequency crushing, analog tape with its natural limitations and digital with its superimposed outboard limiters? As Stripp says himself, "Many defects are far easier to identify aurally than to measure, and they would never show up in exhaustive measurements unless pinpointed in a listening test." In other words, specification comparisons between machines are essentially not 100% comparisons between sound quality; moreover, a myriad of unspecified sonic problems may appear with new technology which can only be truly judged not through pure objective testing but through critical subjective listening. Because of this quagmire of insecurity over digital fidelity, recording engineers have erred on the side of caution by sticking with their well known digital tape machines.

Such a thing as "high-frequency crushing" can be solved by digital algorithms. How is 24-bit really so different from its 16-bit counterpart? To answer this question, some simple math needs to be considered. To get and idea of the vast increase of information resolution between a 16-bit and 24-bit encoder, one only has to compare the available combinations between a binary 16-digit number and a binary 24-digit number: in 16-bit systems, 65,536 levels are possible whereas a full 16,777,216 levels are possible with the 24-bit system. The difference is staggering. Obviously, much greater sonic detail can be achieved with 24-bit quantization.

Related to the increase in resolution power of 24-bit recording is the decrease of inherent noise in the system. Noise is necessarily added to any digital system in the form of dither to blur and encode the discrete steps between quantization levels. Since more quantization levels exist in a 24-bit system, the spacing between each step is thus smaller than exists in a 16-bit system. As the amplitude of dither is calculated to be slightly greater than one quantization step, a 24-bit system will need a smaller dither level to mask the system's digital distortion. Less dither introduced into the digital system equates to less noise and therefore an increased signal-to-noise ratio (SNR). As a rule of thumb, signal-to-noise ratios in digital systems can be approximated by using a 6 decibel per bit of resolution conversion factor. Thus, a 24-bit system has approximately a 148 decibel SNR compared to a 96 decibel SNR for 16-bit reproduction.

Another technical factor which argues against 16-bit master recording and for 24-bit resolution on important musical projects is the nature of digital calibration. With analog tape machines, there exists what is known as headroom above the zero reference level. Typically, this headroom provides up to 18 dB of leeway before the recorded signal starts to become noticeably saturated. With such transient signals such as drum attacks or vocal sibilance, this 18 dB gives an engineer the opportunity to keep most of the musical material around the zero reference level while the occasional peak material falls within the recorder's headroom. Digital calibration, however, is a whole other beast. Digital recorders do not have the leeway of headroom. When signal on a digital recorder exceeds the zero reference level (thus exceeding the quantization resolution), the machine is unable to code the signal and thus distorts (commonly known as clipping). Compared to the smooth saturation qualities of analog tape, digital clipping happens immediately once the signal is too high, resulting in an extremely severe and noticeable burst of noise. Since music naturally has such wildly varying dynamic ranges, though, an artificial headroom must be introduced to digital recorders. Digital recorders are thus calibrated at -18 dB compared to their full signal, thereby allowing for a generally equivalent amount of headroom as their analog counterparts. It is fundamental to the the understanding of SNR in digital recorders that the average program level is set to this -18 dB level. As stated earlier, 18 dB translates to approximately 3 bits of information (1 bit/6dB). A reduction of 18 decibels thus equals a reduction of 3 bits of resolution. Music recorded in a 16-bit system digital tape medium, therefore, only uses an average of 13-bits to encode signals. Ironically, 20-bit resolution is needed to "truly" represent the possibilities of 16-bit sound quality. 24-bit technology is obviously more than a step in the right direction.

Let us turn from the mathematics that goes into comparing 16-bit and 24-bit systems and look to the aesthetic sonic differences between the two formats, for it is this perceived increase in fidelity which is truly the reason for such an upgrade. When the 16-bit format was established, "engineers were aware that this standard severely compromised the quality of the recorded signal," and that, "the human ear is capable of discerning music at sensitivity levels that require 20 bits or more sample resolution." (Levitt) Unfortunately 16-bit was, at the time, "the highest quality that could be attained with the chips and processors available." (Levitt) Said Sony pro audio general manager of new high resolution products, "With these advanced production tools, recording studios are able to benefit from the 24-bit production process that represents a quantum step forward in audio quality." The sonic difference between 16-bit and 24-bit resolution, therefore, apparently represents a "quantum step forward in audio quality."

Despite the almost universal acceptance of 24-bit technology, some proponents of the status quo argue against the need for higher resolution. With a series of listening tests conducted by Mitch Gallagher for Keyboard Magazine, a group of professional musicians and sound engineers were gathered together to compare and contrast master recordings made of 16-bit and 24-bit formats. The end results proved that almost everyone involved could hear a difference between 16-bit and 24-bit resolution. The ironic twist to the test, though, was that some musicians felt the extra fidelity of 24-bit recording added too much clarity to their sound. Apparently, as much as the higher resolution picked up more of the appealing qualities of the sound, it also revealed more of the unappealing qualities of the recording chain, such as microphone preamp noise, string buzz, or room noise. This test brings up an important point about musical aesthetics versus sonic fidelity. Basically, what sounds "good" varies from person to person. Many people, for example, will say one rock band sounds "good" and another "bad" while a different group of people will say just the opposite. Some people think 24-bit sounds "good" and others like 16-bit better. The definition of what sounds "good" therefore is purely a matter of opinion. One thing that can be said about 24-bit audio, however, is that it sounds clearer and more precise to the original sound which it is capturing. Glenn Meadows comments again, "You can hear in the finished CD which one came from the high-resolution mix source. You can hear the improved detail." (De Lancie) It is this increased sonic "detail" which attracts recording engineers to the higher-resolution format.

Another thing learned as a result of these listening tests is that the value of high-resolution audio is only its value compared to other pieces in the audio system. In other words, a chain is only as strong as its weakest link. For example, if less than high-end professional microphones, preamps, and signal processors are used in tandem with 24-bit audio, the resulting signal will probably not sound so much better than its 16-bit counterpart as non-digital devices are limiting the inherent sound quality of the entire audio chain. If a studio does not possess the highest-quality monitoring system, the difference between 16-bit and 24-bit recording may not be able to be discerned. If the engineer cannot hear the difference in quality due to insufficient monitor quality, the engineer will not be able to take advantage of the higher 24-bit resolution while mixing. In other words, 24-bit high-resolution is an all or none process.

The most basic question with high resolution audio, however, is how it will benefit the end consumer who listens to music at home. Why bother with 24-bits if it is only going to be squashed to a 16-bit anyway? The best place to go for the answer to this question is obviously the mastering engineer, for he/she is the one concerned with the transfer of a high quality master to the CD standard. Apparently, the main reason for using 24-bit machines throughout the recording process is to maintain the highest possible integrity of the audio signal until the very last moment when it has to be transferred to compact disc. Quantization noise, much like the noise introduced into an analog system, has the tendency to accumulate with each successive generation. Mastering engineer Glenn Meadows reasons that, "every time we redither or truncate....we limit the resolution, and it compounds down the line." (De Lancie) Another mastering engineer, Ted Jensen, puts it another way: "As you go stepping through the [audio] chain using 16 bits each time, you rapidly start losing dynamic range and resolution and other nice qualities of the sound." (De Lancie) Mr. Jensen does have a point. Digital recorders are not free from inherent noise in the processing of audio signals. Just look at the specs of any digital recorder (a simple DAT machine, for example) and one will find a clearly printed signal to noise ratio. If one stays completely within the digital domain, this noise is prevented from building up. But whenever a transfer to the analog domain occurs (as so often happens during mixing these days), an increase in the quantization noise results from the dithering DAC and the following re-quantization. Basically, therefore, 24-bit allows the highest possible audio chain up until the final point of transfer between the master recording and the consumer.

Even if the master recordings benefit from 24-bit resolution, can the consumer appreciate the increased fidelity of these high resolution masters? Is the use of 24-bit mixdown and recording devices solely for the sake of posterity when presumably consumer 24-bit listening formats will be available? Thankfully, the answer is that much of the increased sonic fidelity of 24-bit recording can be preserved even on a 16-bit Compact Disc through the use of specifically engineered encoding systems. The three currently most popular such methods are Super Bit Mapping, High Definition Compatible Digital, and Apogee's UV22 process.

Probably the more common of the three systems is Sony's Super Bit Mapping. At the recording studio in which I work (Greene Street Recording), both control rooms are equipped with Sony SBM DAT recorders. These high resolution transfer processes, as you may have guessed, are useful not only in capturing the increased fidelity of 20- or 24-bit recording onto the 16-bit format, but also for preserving as much of the sonic information from a high quality 2" master tape. One of the main attractions of SBM to engineers is that the process is compatible with every CD player or DAT machine. No special decoder is required to unravel the high resolution audio. The method by which Sony preserves this high resolution information in only 16-bits is through two common digital audio processes: noise shaping and psychoacoustic principles. By observing the Fletcher-Munson equal loudness contours, one can see that the ear is more sensitive to frequencies in the 500-5,000 Hz range (roughly the range of the human voice). According to Sony, "if the resolution of this 'highly audible at low volumes' zone is improved, the perceived fidelity will greatly increase." Super Bit Mapping therefore shifts quantization noise out of the 3-5 kHz range and redistributes it above the 15 kHz range. This noise redistribution keeps the total noise in the system the same, but reshapes the noise pattern. The underlining of the word "perceived" was added by me because the use of this word brings up a very important point. The true measurable fidelity of a 16-bit SBM source is no better than a standard 16-bit recording. However, the listener perceives the SBM source to be higher fidelity. It is this perception which is the result of psychoacoustic principles. Super Bit Mapping, in fact, does not even accurately follow the hallowed Fletcher-Munson curves because those tests were conducted with sine waves, not music. Sony has apparently done its own testing to determine how listeners perceive increased sonic resolution.

A second system common to Compact Disc mastering is Pacific Microsonic's High Defition Compatible Digital (HDCD). The HDCD system uses a very different approach than SBM. Invented by Keith Johnson and Michael Plaumer in Berkeley, CA, High Definition Compatible Digital necessitates the use of an HDCD decoder on the listener's end to appreciate the enhanced audio fidelity. Basically, high resolution audio which undergoes HDCD processing is first sampled at an extremely high rate (along the lines of at least 200 kHz). A standard 16-bit CD signal is then output from the HDCD processer. Another band of PCM data, however, is also output and recorded as a side band along the standard 16-bit information. When playing an HDCD disc in a regular CD player, only the regular 16-bit PCM code is read while the side band of HDCD encoding is ignored. The true power of an HDCD disc, though, is realized when it is played back on an HDCD compatible compact disc player. The extra side band of digital information is used to faithfully reproduce the original high resolution source. The main disadvantage of HDCD, as compared to Super Bit Mapping or the UV22 process, is that this special decoding system is needed to truly appreciate the benefits of the high resolution audio. Even though HDCD recordings are compatible with standard CD players, large companies are currently not interested in using the system since it requires a increase in manufacturing cost while only catering to the minutely small percentage of the population who currently own HDCD decoding systems. HDCD is therefore caught in a sort of catch-22 situation: Manufacturers will not produce HDCD discs because not enough people own decoders, while people will not bother buying decoders since not enough HDCD discs are being manufactured. It is a shame that HDCD is currently caught in this loop because the sonic fidelity that the process affords the listener is apparently better than both SBM or UV22. HDCD "returns warmth, character, ambiance, and depth to music," which was previously unavailable through a digital medium. (Levitt)

The final mastering process currently in great use among mastering engineers is Apogee's UV22 system. Popular opinion among mastering engineers seems to favor the UV22 as the main high resolution encoding system for compact discs. Ted Jensen of Sterling Sound comments that, "UV22 kept the 24-bit signal perfectly clean...all the way down to -120 dB." I remark that mastering engineers prefer this system in general because the encoding method often depends on the needs of the project. Much in the way that a recording engineer may pass over the favored Neumann U87 microphone to use an AKG

414 to record vocals in a certain situation, many mastering engineers find other encoding means better suited for particular recordings. In general, however, the UV22 system is preferred. Another benefit to engineers of the Apogee system is its general ease of use. The UV22 encoder merely needs to be placed in the signal path between the source and the 16-bit target.

The reason for this preference is that the UV22 method preserves the sound stage and tonal balance of the original high resolution recording. In order to do this, UV22 places "the algorithmically generated 'clump' of energy around 22 kHz" (product literature). This high frequency emphasis is similar to the bias of a magnetic tape recorder. In fact, this analog modeling is used as a selling point by Apogee. The UV22 supposedly adds white noise similar to that of analog tape. This use of analog modeled noise differs from noise shaping and the Super Bit Mapping method which trade a reduced noise floor for a large noise boost at high frequencies. The UV22 system places this boost outside of the audible range. As opposed to dither which adds noise to a digital system, UV22 keeps the audible noise floor (20 Hz-20 kHz) solidly at the theoretical minimum of -96 dB for 16 bits (remember 6 dB of signal to noise ratio per digital audio bit). In order to facilitate this completely flat noise floor, UV22 does not modify the noise floor, but merely makes it transparent, up to 30 dB into it. For example, a 1 kHz tone at -108 dB will distort if truncated at 16 bit information, but will be audible without distortion up to 20 kHz using Apogee's system.

The main reason for this 24-bit revolution is to allow greater distribution and easier attainment of high fidelity music. The appearance of 24-bit digital recorders on the market just a year or two ago is making 16-bit recording devices quickly become out of date and thus pushing down the price of 16-bit recorders which were once the highest attainable level of digital sound quality. Through specifically designed encoding processes such as Super Bit Mapping, High Definition Compatible Digital, and the UV22 System, the fidelity of a high resolution 24-bit source can be appreciated even on a consumer 16-bit format such as the Compact Disc. The new technology of high resolution audio will therefore have and already has had a perceptually noticeable increase in the sonic quality available to the average consumer even though the only people who currently come into direct contact with 24-bit formats are recording engineers themselves. When one considers the usually quick turnover of technology in today's ultra modern society, the two decades of 16-bit digital recording attest to the importance of the current change in the bit resolution standard. Within only a few years, 24-bit audio will most probably be completely pervasive in the audio community whereas the technology of 16-bit recording will be considered obsolete. Definitely, we are seeing just the beginning of a radical resolution revolution in the music industry.

References

Apogee--Product Literature: 3145 Donald Douglas Loop South Santa Monica, CA 90405

Bloom, P. J. "High Quality Digital Audio in the Entertainment Industry." *IEEE ASSP Mag.,* vol. 2, pp. 2-25 (1985)

Craven, Peter. "Toward the 24-bit DAC." *JAES* . V. 41, May '93. pp. 291-313

De Lancie, Philip. "A Bit Too Far?" Mix. 20:12. December 1996. pg. 62

Fleischmann, Mark. "High Definition Compatible Digital." Rolling Stone. 30 May 1996.

Fostle, D.W. "19 bits in a 16-bit sack." Audio. 80:3. March 1996

Fostle, D. W. "Digital Deliverance." Audio. April 1996. 80:4. pg. 26

Gallagher, Mitch. "Seven Things You Should Know About High-Definition Audio." Keyboard. Vol. 24, No. 6, Issue 266. June 1998. pg. 44

Gallagher, Mitch. "Can You Hear the Difference?" Keyboard. Vol. 24, No. 6, Issue 266. June 1998. pg. 48

http://www.tascam.com/news/DA-45HR_PR_V6.htm

Koepnick, Bill and Hodson, Jim. “Format Tales.” Mix Magazine. 21:12 December 1997. pg. 127.

Levitt, Jason. "Alternative Compact Disc Systems." NYU Master's Thesis. May, 1995

Mac, Jeff. "High-Definition Audio." Keyboard. Vol. 24, No. 6, Issue 266. June 1998. pg. 34

Moorer, J. A. "Digital Coding of High-Quality Musical Sound." *JAES* , vol. 27, no. 9, pp. 657-666 (Sept. 1979).

Palenchar, Joseph. "HDCD Catching On At Audio's High End." TWICE. Nov. 28, 1994. pg. 20

Petersen, George. "Digital Consoles." Mix Magazine. 21:11 November 1997. pg. 32.

Petersen, George, "Digital Multitrack Recorders." Mix Magazine. 21:8 August 1997. pg. 106.

Pohlmann, Ken C. Principles of Digital Audio. McGraw-Hill, Inc.: New York, NY. 1995

Rumsey, Francis. Digital Audio Operations. Focal Press: London. 1991

Sony Product Literature: 550 Madison Avenue, New York, New York 10020

Strawn, John, Ed. Digital Audio Engineering. William Kaufmann, Inc.: Los Altos, CA. 1985

Stripp, D. "Digital Audio--A Decade of on Air Operation." Digital Audio. Collected Papers from the AES Premiere Conference, Rye, New York, 1982, June 3-6.

Vail, Mark. "Hardware to Make It Happen." Keyboard. Vol. 24, No. 6, Issue 266. June 1998. pg. 50

Watkinson, John. The Art of Digital Audio. Focal Press: London. 1994

Wilson, Rob. "New CD Recording Process Adds Musical Depth, Detail." The Financial Post. June 26, 1993 s.13

Woram, John M. and Kefauver, Alan P. The New Recording Studio Handbook. ELAR Publishing Company, Inc.: Commack, NY. 1989