The 24-bit Field Recording FAQ

Compiled, written, and edited by Dan Heend and Marc Nutter

Summary
The purpose of this document is to provide information on the 24-bit field recording experience. Topics include all aspects of 24-bit recording, playback, editing, storage management, and archival and trading methodologies. It is based both on research and field experiences of the authors and contributors. Some of the material may imply previous knowledge and understanding of digital audio in the 16-bit realm. While the contents of this document are specifically targeted to the needs and concerns of field recordists, some of the content can be applied to home recording as well. The submission of additions, corrections, and comments, is requested and encouraged, as they recognize that things change rapidly as new products, solutions, and techniques are discovered and released. It is also noted that the scope of the document at inception is markedly heavy in computer-related content, both due to the professional background of the primary author, and the fact that the majority of 24-bit field recording pioneering efforts to date have had the involvement of computers at some stage in the life cycle of a recording. The primary audience is a current digital audio field recordist and audiophile with appreciation of exceptional sound quality and musical content equally. Intermediate level computer skill is a plus, but is not necessarily required for the reader to achieve benefits from this document. Finally, sharing this document with everyone is encouraged. However, the authors request that it is never copied, but rather linked to. Please direct any interested parties to http://www.sonicsense.com/24bitfaq.html so that the most current version of the document can be assured.

We are aware and enthusiastic that many recordists are now seeking alternatives to tape-based equipment and are utilizing laptops with varying degrees of success. As the use of laptops is relatively new, and comparatively experimental in contrast to other recording technologies which have been utilized for many years, we welcome any insights, corrections, or commentary that any reader wishes to share. As high-resolution field recording is in its infancy, we appreciate any contributions that can benefit of field recordists. We will do our best to incorporate such feedback into future revisions provided that it is delivered to us with supporting documentation or first-hand experience as appropriate. While this document has not been prepared as an advertisement for any retailer, manufacturer or product, references are made to currently available products, their feature sets, and to the companies involved in delivering 24-bit recording equipment to recordists. The use of these products is entirely at the discretion of those wishing to implement it. No retailer, manufacturer, or author of this document can take any responsibility for a recordists success or failure as based on the information contained herein. While we share this information in hopes of improving the way audio is recorded, we still suggest contacting a pro audio retailer and computer experts when attempting to set up a computer-based field recording system.

In developing this FAQ, the intent is to share our information and insights to assist all ambitious field recordists and to inspire further discussion and information sharing on this exciting development. Some of the systems and products identified herein are known to have worked well in the field in the past, but as we all continually seek the best possible solutions, we remain open to suggestions that can improve everyone's recording experience. We are all in it together.

Happy Recording!!!

necessarily reflect the views or opinions of the entities hosting it. The primary author reserves the right to edit all submissions for inclusion in this document.

About the Authors:

Written my MJ Parker.

Dan Heend is a 17 year veteran of the computer industry, specializing in telecommunications and network infrastructure engineering, implementation, and management. Outside of work, Dan likes to play an Alembic 5-string bass in his band Next Stop Pluto, and in what little time he has left, he enjoys high quality digital audio recording and mastering, as well as being tortured by his two cats. In April 2000, unsatisfied with the sonic compromise of 16-bit CD quality PCM encoding, he began pushing the envelope with 24-bit recording in the field with an employer-supplied IBM Thinkpad 600E and a Digigram VXPocket. In February 2001, after replacing the Thinkpad with a Sony Vaio Picturebook C1VN, he sold his HHB PortaDAT, satisfied that 24-bit recording was not only practical enough, but also convenient enough to record a touring band performing live night after night. You may have seen Dan running a laptop in the tapers section at numerous String Cheese Incident, Steve Kimock Band, Phil Lesh and Friends, and Funky Meters concerts.

Marc Nutter is a field recordist with 14-years experiences collecting live audio. As a hobbyist, he always sought the latest equipment that his budget could handle at the time. After years of buying, selling, and trading field recording equipment, he started his pro audio company Sonic Sense, which for the last seven years has been dedicated to cutting-edge developments for field recordists of all kinds. Among Marc's numerous projects to push the fieldrecording industry, he has been involved in close relationships with manufacturers to design such products as the Lunatec V2 microphone pre-amp from Grace Design, the MP-2 microphone pre-amp from Sound Devices, the SONIC AD2K+ analog-to-digital converter from Benchmark Media, and the entire line of SONICASE equipment carrying cases. Marc's passion for capturing top quality audio drives him to continue developing new products and means for making the best recordings technology will allow.

In their continued pursuit of leading-edge products, Dan and Marc have used DAARWIN-24, the world's first turnkey 24-bit, DC-powered digital audio workstation, for over 100 hours of successful 24-bit recording.

Revision History

Initial Release June 11, 2001

Table of Contents

1.1.1.1.3 How much RAM do I need?

1.1.1.1.4 How much hard disk space do I need?

1.1.1.1.5 I like the idea of using a laptop, but now I need to be concerned about stability of my recording device. What can I do to safeguard my machine from becoming trashed?

1.1.1.1.6 How can I set my machine up so that it is more stable?

1.1.1.1.7 Can I use the machine for other things while recording?

1.1.1.1.8 My employer gave me a laptop. Can I use it for field recording?

1.1.1.1.9 Are there any other options similar to laptops?

1.1.1.1.10 What brands and models of laptops have been proven performers in the field?

1.1.1.1.11 Should I expect that small, lightweight laptops like the Sony Vaio Picturebook C1VN will get even smaller and cheaper in the future?

1.1.1.1.12 What other features should I look for in a laptop to be used for recording?

1.1.1.1.13 The screen looks small on the Sony Vaio Picturebook C1VN. Shouldn't I look for a laptop with a larger viewable area?

1.1.1.2 Mac

1.1.1.2.1 What models of portable Macs should I consider?

1.1.1.2.2 What are the minimum CPU requirements needed for accurate 24-bit recording?

1.1.1.3 Sound Cards

1.1.1.3.1 What are my options for interface types?

1.1.1.3.2 What types of soundcards are available for portable recording?

1.1.1.3.3 What solutions are available to allow a half or full length PCI card to operate on a laptop for recording?

1.1.1.3.4 If my laptop has two PCMCIA slots, can I use two Digigram VXPocket v2 cards and record 4-track instead of 2-track?

1.1.1.3.5 The Digigram VXPocket v2 card is pretty costly with a streetprice of nearly $600. Do I have to spend this kind of money to buy a 24-bit soundcard for my home machine as well?

1.1.1.3.6 My USB soundcard says it has 24-bit A/D converters, but is only capable of making 16-bit recordings. What’s going on here?

1.1.1.3.7 Can’t I just use the built in sound card on the laptop?

1.1.1.3.8 How can I record more than 2 tracks?

1.1.1.3.9 Everything looks like it’s set up properly. I’ve set my software to record from the digital-in on the soundcard, but it says there’s a problem when I hit record. What’s wrong?

1.3.2.1 If I intend to make 16-bit 44.1kHz CD's from my 24-bit recordings, should I record at 44.1kHz or 48kHz sampling rate?

1.3.2.2 What happens if I set my software to record at 24 bits, but my source is actually only 16 bits?

1.3.2.3 What is 32-bit floating point and how will it affect my recording and editing?

1.3.3 Levels

1.3.3.1 How important is it to get as close to 0 as possible without going over? with 24-bit recording.

1.3.3.2 How can I set my levels to maximize the value of 24-bit recording while continuing to provide 16-bit devices with a good signal?

1.3.3.3 Do I need to set levels my levels on the laptop?

1.3.3.4 Should I use a limiter in the recording process?

1.3.4 Sampling Rate

1.3.4.1 How practical is 96kHz vs. 48kHz currently?

1.3.4.2 My mentality has always been to record at 44.1kHz because I don't want to perform a sample rate conversion. Should I continue to record at 44.1kHz, or am I better off just using 48kHz for the future?

1.3.5 Managing your Power Requirements

1.3.5.1 How can I power all of this equipment in the field?

1.3.5.2 Can't I use a 12V power inverter to supply my devices with 120V?

1.3.5.3 Is there any way I can use lead-acid batteries or gel cells to power my laptop?

1.3.5.4 While recording with a laptop, what can I safely do to conserve power without jeopardizing the recording quality?

1.3.5.5 Which of the power management features should I enable on my laptop?

1.3.5.6 What about the Galaxy Audio Faroutlet as a source of portable power?

1.3.6 Caring for your Equipment

1.3.6.1 How should I store my laptop while traveling?

1.3.6.2 How should I store my laptop while in use?

1.3.6.3 How often should I defragment my dedicated audio partition?

1.3.6.4 My laptop gets pretty hot while in use. Is this bad?

1.3.6.5 What can I do to keep my laptop cooler?

1.3.6.6 It it safe to operate a laptop or hard drive recorder in a humid climate?

1.3.6.7 What do I do if it's raining at an outdoor show?

0.0 Introduction

Ever since the inception of recorded audio, the pursuit of a better, more accurate reproduction of sound has brought mankind from music boxes and player pianos, to Edison Wax Cylinders, through vinyl records, analog magnetic tape, and into the digital age. The journey has taken us from monaural to stereophonic, to quadraphonic and 5.1 surround sound, and finally to the extravagant sound experiences like that found in Sony Imax Theaters, all in attempt to bring to the listener the most transparent musical and acoustical experience of a live source. Kickstarting the digital age for most consumers was the availability of music in a new format, the compact disc. Within a few years, CD's become the primary media of megaplex music superstores. The CD, being a play-only format, left digital recordability relegated to the professional arena with DAT recorders and MD devices.

Before CD's came into commonplace existence, a technical decision was made to determine just exactly how much music would fit on a disc, and what the resolution would be. And in the end, a compromise was made, sacrificing some level of quality in order to make a music media format allowing between 74 and 80 minutes of music to be distributed in a single, easy to use format. While production studios have almost always recorded and worked with higher quality resolution format and media (digital open reel, 30 i.p.s. analog reel, etc.), until recently, such technology has remained far out of the reach of practicality and affordability of the average field recordist.

Today, music recording enthusiasts with a moderate level of computer knowledge and a decent connection to the internet (not all that uncommon anymore) can participate in a whole new world available to the audio recording and trading community. Of course, the quest of the perfect sound is never complete for the recording enthusiast and audiophile alike. With the recent availability of both small and powerful computers, and the flood of new products based on DVD technology, we are left to wonder what the next step will be in the progression of audio technology. Still, as major audio product manufacturers and record labels continue to argue and jockey indefinitely for position as the owner of rights to the new standard of 24-bit media format, field recordists anxious to take advantage of the new technology are left with even more waiting in the wings.

In the past, it hasn't exactly been seen a high priority by audio product manufacturers (some of whom also happen to be music production companies) to make available to the consumer the ability to make perfect clones of digital audio. The implementation of the Serial Copy Management System (SCMS) and associated fees that must be paid for "pro" decks that do not enforce SCMS is a perfect example of this. Additionally, because of the field-recordists' relatively minute market share in the minds of major manufacturers, production of products like the HHB PortaDAT and the Apogee AD1000 are constantly limited or discontinued due to lack of sales and profitability. These companies often cannot justify the development and production costs of niche market products that sell in quantities of less than 100 units per month.

Even with the growing popularity of concert recording, hobbyist nature recording, and the ever expanding art of sound effects and ambience collecting for films and video games, developing high-resolution equipment for these markets, unfortunately, remains a low-priority to high-profile manufacturers. To date, manufacturers have primarily relegated their design efforts, and consequently the benefits of 24-bit digital recording devices, to the controlled environments of recording studios, where A/C power is abundant, professional interfaces like T/DIF and ADAT Lightpipe are commonplace, budgets are usually less of an issue, and product development is more valuable.

Still, we remain confident that 24-bit audio playback solutions and standards will eventually surface in a more convenient format. Therefore, pioneering efforts are being made to ensure that higher quality source material will be available for mastering to new media when the time comes. Just as PCM encoders and video tape machines were once employed against all convenience to take us to the next level of sonic improvements, creative, custom solutions must be employed to capture and store 24-bit audio. Hopefully this document will provide necessary information to help the field recordist realize that it is possible to achieve better sound quality today, rather than waiting indefinitely for marketing interests and technology to evolve. It should be noted that due to the nature of digital recording, coupled with the widespread proliferation of computers, the internet, and the new marketing models they enable, it is the author's opinion that the involvement of the computer-based digital audio workstation in the life cycle of 24-bit recording(field or otherwise) will probably never subside. This FAQ was written with the uncompromising yet realistic field recordist in mind.

0.1 Background

0.1.1 Why 24 bits? Isn't 16 bits enough?

A: 16-bit technology is over 20 years old, but that's not simply a reason in and of itself. The Compact Disc format and specification that promised superior quality audio was derived from a compromise between playback time and sonic quality. It is this compromise that facilitated putting 16-bit audio within reach of everyone's home, car, and person. However, it is this same compromise that has left audiophiles to grasp tightly onto their old vinyl LP's and their audiophile-grade turntables in hopes that something better would come along. Today, 24-bit audio technologies promise to deliver the same sonic experience once promised by CD's. The notable difference here is, 24-bit technology actually delivers the level of performance most audiophiles expected long ago. Sure, nothing beats the live performance experience, and 30 i.p.s. 2" analog reel tape is second to none in its ability to reproduce that experience. But we can surely say that the edge 24-bit recording has over their 16-bit counterparts puts a big smile on the audiophiles face. Why? What's missing on my 16-recording?

Simply, the answer is detail. The PCM format provides its optimal resolution when signal levels are at their very highest. As signal levels decrease to lower levels, resolution deteriorates, leaving quiet cymbals and string instruments sounding typically sterile, dry, harsh, and lifeless. The more bits you have available to you in the process of quantizing the amplitude of a waveform at any given sampling, the more accurately a lower level signal can be represented. If an instrument is very loud while standing next to it, but is recorded at a low level, there are less numbers that can be used to represent just exactly how loud it is at any given moment. We know that a wave modulates between silence and its maximum amplitude or volume, while the number of times per second this modulation occurs gives us the pitch of the wave.

For example, if there are only 4 discrete numbers that can be used to represent the volume level of a particular recording, 1 would be silent, 2 would be very audible, 3 would be louder, and 4 would represent the loudest level. Can you imagine what all your audio would sound like if these were the only choices for representing amplitude at any given moment? Definitely horrible, and it would sound like square wave distortion and noise. This example would be a 2-bit recording.

In order to make this sound better, we need to be able to have discrete values in between these values. A fading piece of music can't just go from very audible to silent, or it wouldn't be a smooth fade. A 4-bit recording would have 16 discrete possible amplitude levels. Can you again imagine what this would sound like? Definitely better, but its still a totally horrible representation of the sound. We can deduce from this that the more discrete values available to us, the better it will sound. Is there a limit to the human ear's ability to perceive these inaccuracies? Definitely, but it unfortunately does not stop at the 65536 discrete values afforded to us by 16-bit technology.

The overriding concept here is called dynamic range, and is measured in dB. The dynamic range of a recording is the difference between its loudest point and its quietest point.

To elaborate further, each bit gives us the ability to represent about 6dB of dynamic range. A passage that is 6dB louder than another passage is said to be twice as loud as the other passage. In the 4-bit example, we theoretically have 24dB of dynamic range that can be used. But what if recording doesn't take advantage of all that dynamic range? What if the recording never peaks beyond 6dB of its maximum possible limit? In this case, the recording would only take advantage of 3 of what we call the least significant (or left-most) bits, meaning 18dB of dynamic range. 16-bit recordings are capable of a theoretical maximum limit of 96dB of dynamic range. This means that a single wave could have up to 65536 discrete values that can be used to represent it. But if the same wave recorded at 16-bit peaks at 48dB below its maximum possible limit, then there would only be 256 discrete values that can be used to represent it, taking advantage of only 8 of the least significant bits. The 8 most significant bits would contain no information whatsoever, and would remain unused. In the case of 24-bit recording, you'd have a maximum of 16,777,216 values to choose from, and in the case of a wave peaking at 48dB below its maximum possible limit, the wave would still have 65536 possible discrete amplitude values that could be used to represent it.

Now, have you ever heard any of the early 8-bit computer recordings that floated around in the early days of home computers? Didn't they sound just awful? I mean, you were impressed because you had a snippet of music that you could recognize playing from your computer, but you wouldn't want to listen to it for more than a minute or two. I personally remember playing back an 8-bit digitized 5 second snippet of Van Halen's rendition of the Kinks' "You Really Got Me" over and over again on my Atari 800 until I couldn't take it anymore. The thrill soon had me building an 8-bit digitizing device with a microphone input jack and a connector for the joystick port. Ah, those were the daysà but I digress.

Perhaps many are more familiar with 8-bit audio from real-time internet sources like RealAudio. It's good enough for speech recognition, but leaves all too much to be desired for music.

Now here's the kicker in the16-bit realm. While the volume level of a recorded low-E note struck on an acoustic guitar might take advantage all 16 available bits (for instance, where the peak on the DAT deck reaches 0dB), the squeak of the fingers on the string, the scratch of the pick hitting the string, and the 5 or 10 audible harmonic overtones of that note may never reach a point beyond 48dB shy of the 96dB maximum. Yes, all of these additional by-products of that low-E string that make the guitar sound alive and compelling receive all of the fidelity of that scratchy, distorted, computerized sound of that 8-bit sample from long ago. And as the basic low-E note fades out, it too gets the same butcher treatment from the ever decreasing number of discrete amplitude values. Yikes!

Now record with a 24-bit word length, and put the CD quality back into those string squeaks, pick scratches, and overtones. With 24 bits, you can hear the clarity of the cymbals decaying as they keep ringing smoothly down to complete silence. The little low-level smack of the bass pedal head hitting the bass drum skin that sounded barely like a small click before (if audible at all) now sounds like a smack, complete with its own smoothly reverberating decay. Even the low-level acoustical reflection from the wall behind the band now contributes to the experience with added detail and a sense of ambience, not simply low-level distortion. Finally, because of this improvement, no more does the recordist have to risk overloading and clipping the recording in effort to achieve maximum fidelity. Levels can be set conservatively with the assurance that a high degree of fidelity is maintained.

0.1.2 What does the future hold for 16-bit audio recording devices like DAT machines?

A: While the DAT recorder will certainly still be around for a while, Alps (the company that makes all of the DAT transports) has recently stopped production of 4mm tape transports for audio use. They will continue to manufacture transports for computer backup devices, but once the audio transports are gone, that will mark the end of DAT machine production. Tascam has verified this fact, and has followed by stating that it has enough transports for another 2-3 years of production. Therefore, more and more companies (like HHB) will be phasing out their DAT machine production in the near future. Word has it that Sony also recently stopped production of the D7 and D8 decks (rumor?). They will continue to make the M1 and D100, but for how much longer? One must seriously question whether good service and parts will be readily available tomorrow for a DAT deck bought today, especially in light of the recent proliferation of 24-bit capable recording and playback solutions. Additionally, as long-term DAT recordists know all too well, DAT tape is far from long-term permanent storage solution. In most cases, a five-year-old DAT tape will show marked deterioration in sound quality as the effects of electro-magnetic fields, friction, humidity, and temperature take their toll. Therefore, people looking to enter the digital audio recording arena today should take a very serious look at 24-bit technologies if they want to avoid the inevitable obsolescence of 16-bit audio formats.

0.1.3 What does the future hold for the playback of 24-bit audio?

A: Stand-alone devices capable of DVD-Audio playback are now becoming available. SACD, which promises performance similar (or arguably better) to 24 bit resolution, is another new format being promoted by Sony/Phillips, and music titles for both DVD-Audio and SACD formats are slowly becoming more available. Recordability for DVD-Audio is on the horizon as well for consumers, while SACD (DSD) recordability is currently only available to production companies willing to pay the huge price. The DVD-Audio format is said to currently have a leg up, with more record production companies on board, and due to the fact that DVD players have replaced VCR's in new sales and in many homes. Several manufacturers are also beginning to release DVD players capable of reading DVD-Video, DVD-Audio, as well as standard CD's, and even SACD's.

0.1.4 If there isn't an easy format for storing and playing 24-bit audio, why should I bother recording in 24-bit?

A: In addition to the basic sonic improvements 24-bit recordings provide over their 16-bit counterparts, convenient storage and playback formats are in fact just around the corner. In the past, the proliferation of a better quality playback device has always been very dependent on the availability of recordings in that format. A good example of this is the BetaMax video tape format. We know that 16-bit audio has been part of the past 20 or so years, and there is an undeniable push to go beyond that technology from every facet of the entertainment and audio industry. The fact is, 24-bit recording and storing is possible and viable today. Even if playback devices aren't in every home, the more 24-bit recordings being made today (stored in any format), the sooner we will have a significant library of material for the new format of tomorrow. Not being able to readily distribute 24-bit material to the average person today is a very temporary condition, and should not be seen as a deterrent to 24-bit recording where it is easily possible. In the future, it will be very easy to convert 24-bit recordings on temporary storage to the more ubiquitous standard playback format. Because of low demand, it isn't likely in the near future that portable 24-bit recording product solutions targeted to the average consumer will be available, so we might as well start today with what we have. Besides, it just sounds better!

0.2 Theory

0.2.1 Resolution

0.2.1.1 What is this concept of bits?

A: See 0.1.1.

0.2.1.2 How much improvement do these extra bits make to the quality of a recording?

A: See 0.1.1.

0.2.2 Dither and Noise Shaping

0.2.2.1 What is dither?

A: Dither is the general name used to describe the noise that must be applied (or the process of applying noise) to reduce the perception of digital artifacts during word length reduction operations (i.e. Bit depth from 24- down to 20- or 16-bits) for compatibility with lower resolution devices. In the case of DAT recorders or CD's, which only support 16-bits, dither should be applied to recordings that are mastered or edited in 20- bit or higher systems. The failure to properly dither when reducing word length is referred to as simple truncation, which merely omits some number of least significant bits of each sample. Word length reduction of any kind will introduce some level of PCM quantization noise or artifacts into the signal, resulting in harmonic distortion, particularly affecting lower level signals. Dither can mask these artifacts by adding a small amount of noise to the signal, allowing the listener to better perceive low-level sonic information below the noise floor introduced by quantization.

The most basic of these processes is called Triangular Probability Distribution Function (TPDF) and involves the placement of broadband white noise across the audio spectrum. There are varying degrees of dither, including ½-bit, 1-bit, and 2 bit depth. Use of dither may or may not be accompanied by a type of noise shaping. It is generally recognized that dither noise by itself can be very audible as compared to any other form of word length reduction that also employs some level of noise shaping, because the frequency spectrum of dither noise is white, or evenly distributed across the entire spectrum.

0.2.2.2 What is noise shaping?

A: As all word-length reduction operations utilizing dithering requires the addition of some level of noise, noise shaping is usually (but not always) employed as well, the term "noise shaping" is occasionally used to loosely describe the process of both dither and noise shaping during word-length reduction. However, it should be noted that dither and noise shaping are two very different things. In the case of the Benchmark's dither and noise shaping options on their AD2K+ A/D converter, their use of the term identifies their proprietary algorithms that expand on near Nyquist formulas and the latest research from the dither experts at the University of Waterloo, Ontario, Canada.

Specifically, the goal of each noise shaping technique is to shift the necessary dither noise into frequencies where it will be least perceptible during playback. The choice of noise shaping used depends on several factors, including personal preference, the stage in the life cycle of the recording, and the type musical content. For instance, a particular noise shaping curve may place dither noise in higher frequencies, which may be desirable when recording a kick drum, where it is known that the kick drum will later be equalized to reduce these high frequencies (thereby reducing dither noise). Increasing noise, or "energy" to high-frequency sounds like cymbals can have the effect of making the recording sound unnaturally harsh or exaggerated in the high-end, or unbalanced.

Some noise shaping algorithms, like Apogee's UV22 employ a constantly changing curve, in an attempt to keep noise in frequency spectrums least likely to be heard at any given moment because of masking by changing musical content. Different noise shaping algorithms have different coloration effects on audio.

0.2.2.3 Do I really need dither and noise shaping?

A: Any time a word-length reduction is required, some amount of dither and/or noise shaping should be employed. Word-length reductions are required in many stages of the life cycle of the recording. If an A/D converter samples audio at 24 bit, but the storage medium used to record the samples only supports 20 or 16 bits, there will be a word-length reduction. When running digital processes on a wave file, extra bits are usually required to store the resultant accurately. However, where the result file must remain at the same word-length as the original file, a word length reduction must occur. If the recordist is utilizing 24-bit equipment but must produce CD's or DAT recordings, again word length reduction is necessary, and dither and noise shaping are essential to minimize the perception of digital artifacts. It should be noted that some levels of dither and noise shaping are designed to be used as a finishing process before printing(e.g. burning a CD, writing to tape, etc.), while others have less impact on the recording and are better suited to being used on audio that may undergo further processing.

0.2.2.4 What types of dither and noise shaping are there?

A: While TPDF is a generally acceptable and very basic dithering algorithm without noise shaping, more complex forms of word-length reduction involve strategic placement of dither noise into frequencies where it is less likely to be audible in an attempt to more closely resemble the higher word length they are striving to emulate. During loud segments of material, dither noise added (shaped or not) during word length reduction will often remain imperceptible. During moments of low volume, the type of noise added may have considerably more audible impact.

Additionally, the placement of the noise can lend itself to improved ease of editing in post production work. In a case where high-frequencies are the target of the recording, using a word reduction that places greater amounts of noise in lower frequencies may be beneficial since they can be omitted in the final mastering process.

Manufacturers of A/D converters like Apogee, Benchmark, and Sony have each refined and packaged word length reduction with dither and noise shaping into trademarked processes like UV22™, NN1™, NN2™, NN3™, NS1™, NS2™, NS3™, and SBM. Apogee's UV22™ and Benchmark's NN™ (Near Nyquist) take into account the same prominent research from the University of Waterloo in Ontario, Canada. Rob Wanamaker and Stanley Lipschitz have made their careers in digital audio with a focus on word-length reduction. Major manufacturers have applied the algorithms that these experts have shared. One basic premise is the attention to the Nyquist limit which provides that human hearing only extends to 20kHz and that any noise placed at or above this frequency will remain inaudible. However, due to frequency range restrictions brought about by sampling rates of 44.1kHz and 48kHz, it is not possible to place all the noise above this point. Consequently, some companies have taken to applying the noise in areas where it is less likely to be perceived or can later be edited out, depending on the type of recording being performed.

Also of note is that dither and noise shaping does have an audible impact on the music aside from the addition of noise. Oddly enough, some of these sonic characteristics may not show up on test equipment but still remain perceivable to the ear. Some will accentuate bottom-end, while others may emphasize mid-range, often adding presence in one area but compromising another.

Word length reduction with dither and noise shaping may also be performed in software. The Waves L1-Ultramaximizer DirectX plug-in provides many choices in both dither and noise shaping employing their IDR proprietary algorithms. Apogee also recently made UV22 available as a software VST Plug-in to be used with Steinberg products like Cubase and Nuendo. The latest versions of Cubase and Nuendo actually come packaged with UV22.

0.3 Listening

0.3.1 I've heard people say that you can't really hear the difference between 16 and 24-bit recordings. Is this true?

A: Not at all. They simply need to turn it up! Seriously though, low-level listening will not allow you to easily hear the benefits of additional resolution of low level signals. The best way to perceive all of the benefits of a 24-bit recording is to listen as loud as possible without damaging your hearing. For concert recordings, this would be considered concert level (about 100-105dB maximum). When most people say it is hard to hear the difference between 16 and 24-bit recordings, they are referring to one of three things. One possibility is they either they don't know what to listen for or how to listen for it. Another possibility is the fact that a high quality playback system is essential, and your average car stereo or boombox just won't cut it. Finally, difficulty in perceiving the difference has also been said in reference to 24-bit recordings that are cut down to 16 bits for CD's before listening, and this is only partially true, as performing processes at 24-bits before converting to 16 bits as a last step before listening does yield a more accurate recording.

The differences to expect are greater realism and more accurate portrayal of the source event. Not so ironically, it may take several listenings, and even a variety of source material before the differences are recognized. We have had the enjoyable opportunity to sit with many listeners and observe their reactions. Sometimes the recording of a rock 'n' roll event, a well-defined bass line, a stunning cymbal shot followed by persistent chimes, or the feeling of being back in the venue that is the kicker that turns them on. For others. it is the dog barking, the door bell ringing or the car we have recorded that really moves them. Either way, with a few listenings, everyone soon hears the impact that 24-bit recording delivers over its 16-bit counterpart.

0.3.2 What do I need to do to set up my stereo to be able to hear the difference?

A: ***Coming soon.

0.3.3 What is the best way to compare 16-bit material to 24-bit material?

A: Recently, the folks at Sonic Sense have been making a number 16-bit and 24-bit comparative recordings. In each case, a field recording system has been set up to include the Benchmark Sonic AD2K+ since it provides simultaneous 16-bit and 24-bit output which can respectively be sent to a DAT machine and the laptop. The 16-bit recordings can then be sent into a computer beside the 24-bit recordings and an A/B comparison can be done with ease.

Alternately, a 24-bit recording can be used for playback through the digital input of the AD2K+ in 24-bit mode or any of the various word-length settings. By turning the front panel dial, the listener can hear the 24-bit feed or any of the 16- or 20- bit outputs. Once changed to any of the 16-bit outputs, the loss in low level detail (particularly in the high-end), depth, and definition of all frequencies will be apparent.

Both of the above comparisons have been made using the premise that a 16-bit feed (24-bit reduced to 16) from the Benchmark converter is superior to the 16-bit conversion available on any DAT recorder.

These sample recordings are available on a complimentary CD storing .wav files for anyone who is interested. To analyze them, one must have a 24-bit soundcard and software that supports at least 24-bit word length. High-quality, external D/A conversion will reveal the differences even more thoroughly than a typical soundcard in a workstation.

0.3.4 There is often a lot of ambient noise recorded in a live concert situation. Won't this minimize the benefits of the additional 8 bits of resolution?

A: No. Even though there may be a lot of crowd noise that may cause audio masking of some lower level signals, other lower level signals that reside in frequency ranges far away from that of the crowd noise may not be masked. Additionally, because ambient noise is random in nature of both frequency and volume at any given point, lower level sounds that remain constant can be heard through the ambient noise.

0.3.5 Which is said to provide greater benefit to the listener over conventional 16-bit 44.1/48kHz recordings: Increased sampling rate or increased word length?

A: Increased word length. From Ken Pohlmann's "Principles of Digital Audio," 4th edition, p390-391:

"The use of high sampling rates such as 96 and 192kHz may seem unnecessary. In rare cases, a person may be able to hear frequencies of 24 or 26kHz, far below the cutoff frequencies of 48 and 96kHz [for 96 and 192kHz sampling rates respectively]. In most cases, high frequency hearing response is below 20kHz. Thus, for steady-state tones, the higher frequency response may not be useful. However, it can be argued that high sampling frequencies improve binaural time response, leading to improved imaging. For example, if short pulses are applied to each ear, a 15ms difference between the pulses can be haerd, and that time difference is shorter than the time between two samples at 48kHz. Some people can hear a 5ms difference, and that corresponds to the time difference between two samples at 192kHz. In theory, this high sampling rate may improve spatial imaging. Thus, it may take two ears to distinguish between a recording at 48kHz, and one at 192kHz."

Of course, this additional sonic information becomes more useful in a recording where a large amount of stereophonic information is present at the source, such as an onstage recording. Less of a difference would be heard between recordings made from further away from the source.

It has also been said that the frequencies typically out of the range of human hearing might contribute to the energy of lower frequencies when the higher frequencies and their harmonic overtones hit the human ear. Still, the increased dynamic range of 24-bit over 16-bit is definitely said to be the more easily perceivable of the two.

It is important to note that the step from 16-bit to 24-bit only requires a 50% increase in storage whereas doubling the sample rate doubles the necessary storage space. While increased sampling rate delivers improved audio quality, the ratio of improvement to file size may not be justifiable when space and time is at a premium. When file size and editing time is an issue, most recordists chose to step up to 24-bit as opposed to going to the higher sampling rate while continuing to sample at 16-bit. Later, as storage becomes less expensive and density of recording media increases, these issues will become less of a factor, and sampling rates like 96kHz will become quite common.

0.3.6 I'm still not sure how much of a sonic improvement 24-bit recording makes over 16 bits. What can I do to better understand the impact the additional resolution has on my recordings?

A: Get some 24-bit sample recordings, read this FAQ. If you have a 24-bit soundcard, try the following. Make a 24-bit recording, keeping peak levels at below -48dBfs. Now make the same recording at the same level, but this time at 16-bits. Now, normalize (peak, not RMS) both recordings so that loudest peak hits 0dB. Finally, run these recordings through an A/B comparison, preferably on the computer, with the listening level turned up as loud as is comfortable. The 16-bit recording will sound very different from the 24-bit recording. You be the judge.

0.3.7 What is a noise floor?

A: The noise floor is the minimum sound present in a recording or given environment. While manufacturers present self-noise specifications on most equipment, this noise is not an issue in a many field recording environments because the self-noise is often much less significant than the typical ambient environment. If you are able to hear your equipment's noise in your recordings, you should consider alternatives in equipment. In the case of a well designed studio incorporating substantial sonic dampening and high quality amplification components, there will be a very low noise floor to the sound source. Conversely, the noise floor in an arena full of enthusiastic concert-goers will be much higher, but will exist in only a limited band of frequencies. This noise may or may not be masked, depending on the frequency and volume content of the recording.

The concept of a noise floor is particularly important when determining the achievable dynamic range of a recording. As defined earlier, dynamic range is the difference between the loudest and the quietest sound of a recording. While many high resolution converters are able to achieve dynamic range in excess of 110dB, the dynamic range of a recording may be substantially less, depending on several factors such as noise and recording levels.

It should be noted that while the noise floor may be higher than that of the dynamic range of a recording, it may still be possible to hear sonically important information below the noise floor. The further down you go below the noise floor, the lower the signal to noise ratio becomes, until the noise eclipses (or masks) the listener's ability to perceive the sonically important information. Dither is used to aid in the perception of sonically important information below the noise floor created by the act of word-length reduction. It effectively raises the noise floor slightly (increasing noise overall), but has a way of increasing audibility of sonic information below the harmonic quantization noise inherent by word-length reduction. Without dither, the ability to perceive low-level information would stop at the noise floor created by the truncation.

0.3.8 Where can I get samples of 24-bit recordings?

A: Look for announcements of 24-bit recording offers on the laptop-tapers@yahoogroups.com mailing list, or other mailing lists like DAT-Heads. Talk to or email known 24-bit recordists and ask for a sample. In most cases, they will be happy to oblige with a concert or sound effect recording.

1.0 Recording

1.1 Hardware

1.1.1 Portable Computer Solutions

1.1.1.1 PC(x86) Architecture

1.1.1.1.1 What's the minimum recommended CPU speed required for error-free 24-bit recording?

A: A 400MHz Intel Pentium II or equivalent is sufficient to keep up with a 24-bit digital audio stream without dropping any samples. While many people have used home machines as slow as a Pentium Pro 200MHz for simply recording DAT's onto the hard drive, higher speed processors are going to prove more stable, especially in a laptop which may be more prone to interrupts like that from a power management event. Since any new machine will have a processor at or above 600MHz, this issue should not be of much concern for new laptop owners. However, the other issues explained in this document are still of considerable concern for maintaining an error-free stream of 24 bit audio into the machine. Even the slightest digression from writing to the hard-drive can cause dropped samples, depending on many factors including the software and operating system used and the architecture of the machine.

1.1.1.1.2 What operating system works best for 24-bit recording?

A: Windows NT 4.0 and Windows 2000 Professional have shown to be very stable, reliable performers when used for digital audio. They are the only Intel based Microsoft OS's that should be considered for recording use. While many sound card manufacturers are striving for compatibility with Windows 95/98/ME, these operating systems are poor platforms for audio recording. The lack of true multitasking in these OS's puts the recordist at high risk of dropping samples while recording. From experience, samples can be lost very easily by merely moving the mouse, closing the display lid, causing power management events (e.g. plugging in an A/C adapter while running) and otherwise changing the focus from one window to another. Program stalls for unknown reasons have been known to occur, sending streaming bits of digital audio data into the abyss. Additionally, because of the architecture of these "home user" OS's, an errant or poorly written driver or program can cause total system instabilities without a moment's notice, leaving the recordist staring at a blue screen. Windows NT and 2000 Professional employ a true multitasking and multithreaded architecture, and programs and drivers are generally well protected from stepping on each other in ways that frequently occur on 95/98/ME machines. Still, the importance of a cleanly installed, properly patched and upgraded OS cannot be understated. The use of something as simple as a mismatched or outdated driver can easily cause system instability. In any case, an OS installation should be thoroughly tested before relying on it in the field.

Users who insist on avoiding Microsoft OS's by using Linux instead will undoubtedly run into the driver (or lack thereof) support issue for at least the soundcard. Yes sir, Billy Gates certainly does have us PC guys by the short hairs.

1.1.1.1.3 How much RAM do I need?

A: Successful recordings have been made regularly on PC laptops with as little as 64MB of RAM. This is probably the least amount of RAM a recordist can safely get away with. WinNT 4.0 and Win2K require at least 32MB to boot, and most recording/editing packages will not need more than 16MB additional to run. 128MB of RAM is sufficient for most software to run with minimal paging or swapping of memory to disk. There is no inherent need to go beyond 128MB of RAM for field recording on a laptop. More RAM might be desired if the laptop is to be used for other tasks besides recording, but 128MB is ample for recording by itself.

1.1.1.1.4 How much hard disk space do I need?

A: An average 3 hour event recorded and stored at 24bit/48kHz requires between 3 and 3.5 GB of hard disk space when utilizing software that operates in standard 24-bit word lengths as opposed to 32-bit floating point. In order to reduce the likelihood of corrupting the OS and software by accidentally filling up your system partition, and to minimize fragmentation, audio data should be written to a dedicated partition, and not to the same partition the OS and software are installed on. Depending on the OS configuration, the amount and type of software installed, and whether you have a dual-boot setup on your laptop,1.5GB - 2.5GB should be allocated for the system partition. Therefore, the practical lower limit of total hard disk size for field recording is about 6GB. A 9GB drive will allow 2 three-hour recording sessions to be recorded easily, and a 12GB drive should enable recording about three performances before needing to worry about offloading the audio files. Keep in mind that when a manufacturer states that a drive is 6GB, they probably mean more like 5.7GB. 9GB drives usually support 8.5GB of usuable space, and 12GB drives yield about 11GB of real space.

Tricks such as the use of lossless compression can help to extend record times a bit, but compressing files initially requires a bit of free space for the output of the compressed file. More about these tricks in the storage management section of the FAQ.

1.1.1.1.5 I like the idea of using a laptop, but now I need to be concerned about stability of my recording device. What can I do to safeguard my machine from becoming trashed?

A: See the DO's and DO NOT's in section 1.1.1.1.5. As a general rule, install as little as possible. It is generally accepted in the computer field that the more you install on a machine, the less stable the machine gets. It is also generally accepted that the longer a machine is running, the less stable it becomes (surprise!). Finally, don't let associates use your machine, unless they are well versed or trained on its proper use. Boot-up and login passwords work very well in protecting your machine from unauthorized use.

1.1.1.1.6 How can I set my machine up so that it is more stable?

A: There are a number of very important rules to remember when setting up your machine from scratch. Install the OS on its own partition, separate from where you plan to store your audio recordings. Software should be installed on the OS partition. Flash your laptop's firmware with the latest revision from the manufacturer's website before installing your OS. Do not install or enable network support in the OS installation you plan to use for recording. Use a dual-booted OS installation for network support and non-recording functions. Build your laptop using the following procedure in strict order: Operating system, basic operating system patches, video driver, and basic laptop support drivers (enabling all ports, PC Card slots, CD reader, and power management) and laptop model specific patches. Then test the machine out very thoroughly. Connect your soundcard and install its drivers and support software freshly downloaded from the manufacturer's website. Test your soundcard with your basic operating system sounds, and then turn all system sounds permanently off, as they can interfere with recording software.

Next, install your audio recording/editing software and test. If you have problems, stop and troubleshoot before continuing. If everything seems to work fine up to this point, go ahead and install other essential operating system patches (e.g. patches to Internet Explorer 5.0, security patches, etc.) Test again to see that the addition of patches did not malignantly affect the function of the basic laptop, nor the ability to record reliably. If you experience problems after this step (not likely, but certainly possible), you will have to start over, and slowly add patches one by one and testing until you find the offending patch.

Finally, you may want to add support for your CD burner, CD burning software, some lossless compression software, and some archival software.

Some DO's for your whole machine:

DO consider a separate installation of the OS for non-recording/editing functions, including network support, word processing, e-mail, and internet browsing.

DO stop PC cards through software before removing them from the laptop.

DO extensive testing of your machine at home before taking it into the field. See what happens when it gets hot, see what happens if you close the screen while recording.

DO disable all non-essential services in your audio recording installation, including items in your system tray, schedulers, and camera support. Save this for your non-recording installation.

DO install virus management in your non-recording installation.

DO consider purchasing a turnkey system from a dealer specialized in designing field-oriented Digital Audio Workstations.

And now for the DO NOT's, specifically for your audio recording installation:

DO NOT install Internet Explorer v5.5.

DO NOT install MS Office (any version).

DO NOT install any additional network or internet software, like FTP downloaders, RealAudio, FlashPlayer, or Shockwave. You can install WinAmp if you like.

DO NOT enable any kind of power management that will power down the PC Card slot, hard drive, or screen after any period of time. Set devices and ports to either always on, or always off. Laptop screens will turn off when the lid is closed or almost closed.

DO NOT enable suspend or hibernation modes of power management.

DO NOT install anything that doesn't have to do with recording, editing, playback, storage management, or archival. Save your games, DVD playing, and built-in camera snapshot software for your dual-booted non-recording installation.

DO NOT use NTFS for Windows NT 4.0 and Windows 2000 installations.

DO NOT install virus protection.

DO NOT allow people unfamiliar with your machine to touch it.

DO NOT use any type total file system compression. This will only slow down operation of your machine as components need to constantly compress/decompress on the fly during use.

If you've followed these rules and procedures when building your machine, you will have a very stable recording platform.

Once again, you may want to seek the assistance of a computer professional or company who can offer these services. While a savvy computer user may be able to set up the system following the suggestions above, the value of system build-out with a warranty cannot be overstated.

1.1.1.1.7 Can I use the machine for other things while recording?

A: Can you? Possibly, depending on how thoroughly you stress tested your machine, the amount of RAM, and how fast and optimized the hard drive, CPU, and operating system is. Should you? Absolutely not.

Asking the machine to do anything else while recording is just asking for problems. Even acts as insignificant as moving the mouse can cause noise and dropped samples to appear in the recording. Anything you do to distract the machine from taking in samples and writing to the hard drive puts the recording at some level of risk. Don't even monitor the levels on the laptop if you can avoid it, and definitely do not enable playback while recording. Stress tests performed successfully at one time at home may not necessarily remain valid for the life of the machine as additional software and plugins get installed and the partitions become fragmented. Resist the temptation to mess with the machine after you hit record, and then close the lid if it doesn't affect the recording (you've tested this, right?). 24-bit recording does help to capture more sonic information, but nothing beats being at the event (and paying attention to the music) for sound quality.

1.1.1.1.8 My employer gave me a laptop. Can I use it for field recording?

A: Sure, but your employer may enforce that your laptop's OS retain the company build. Install your own dual-booted OS for audio recording. As long as it meets the minimum specifications in sections 1.1.1.1.1-3, and you do the proper testing, the best laptops in the field are the free laptops in the field.

1.1.1.1.9 Are there any other options similar to laptops?

A: Good question. A few companies like Dolch (www.dolch.com) make one or two PCI slot "portable" computers. These are usually fully functional desktop PC's, shrunk down into a small box with a flat-panel monitor in it, used to troubleshoot networks. Finding a way to power these devices (outside of 120VAC) though is sure to be a challenge.

1.1.1.1.10 What brands and models of PC-based laptops have been proven performers in the field?

A: Experienced laptop field recordists have had great success with the Sony Vaio Picturebook C1VN ($1899). The Picturebook C1VN employs a recently developed, low power x86 compatible 600MHz CPU called the Transmeta Crusoe, which is responsible for the greatly extended battery life of this unit. See www.transmeta.com for details on this new CPU. It comes with a 12GB drive, 128MB of RAM, 1 PC Card (PCMCIA) slot, 1 USB port, 1 iLink400 port (a powerless, 4 pin version of the IEEE1394 Firewire specification), a Memory Stick slot, a 56kbps modem, and an embedded camera. The screen is 8.9" XGA, allowing 1024x480 to be easily visible, and supporting 1024x768 and 1280x1024 by means of virtual desktop sizing. Battery life with the stock battery is said to be 2.5-5.5hours of operation, and larger capacity batteries are available, including a "quad-capacity" battery easily yielding 8-20 hours of continuous operation. The unit weighs between 2.2-4.5lbs, depending on the battery used, and has the approximate footprint of a Tascam DA-P1 DAT deck. See http://www.sonystyle.com/vaio for more details on the Picturebook C1VN or other potentially viable Sony laptops.

An IBM Thinkpad 600E was used for some of the earliest 24-bit field recordings as well.

This machine has an Intel PII-400 CPU, 64MB RAM, a 9GB hard drive, and a much larger screen. Battery life was a mere 1.25-1.5 hours, so an external battery powering scheme using large 7.2AH or greater, 12VDC lead-acid batteries and a modified 12V laptop car adapter was required to run this rig for longer recording sessions.

We have also received positive commentary from an audio industry professional who is utilizing the HP Omnibook 6000, although they are generally running it from AC power and have the advantage of not having to close the display.

No doubt there are other successful brands and models capable of providing a stable field recording platform. Some newer small IBM Thinkpads look potentially useful according to specs alone. Hopefully, as more recordists begin to experiment with the laptop solution, we will be able to get them documented here.

1.1.1.1.11 Should I expect that small, lightweight laptops like the Sony Vaio Picturebook C1VN will get even smaller and cheaper in the future?

A: To say that we have reached the absolute limit to the minimum size a laptop can possibly achieve would be extremely arguable. However, there are several factors involved which will certainly dictate reality in the practicality of a laptop design. Firstly, basic ergonomics involving keyboard size, layout, and pointing device (mouse) placement has been a major factor in the make or break success of "subnotebook" computers in the past. Products such as the IBM Thinkpad machines with the butterfly keyboard and the Toshiba Libretto sub-notebook with its tiny rectangular keys (overly dense key pitch) and difficult to use pointing device, were all part of machines that sold well initially because of industry hype behind their small size. In the case of the Toshiba product, the product form factor was a long-term failure, becoming prematurely obsolete due at least in part to its ergonomically unfriendly keyboard and mouse. It is the author's opinion as he composes this entire document on a Sony Picturebook C1VN that if the keyboard got any smaller, it would be notably uncomfortable and fatiguing. The keyboard on the Picturebook takes up the entire bottom of the laptop, and fingers rest easily on keys next to each other without being crowded or cramped.

Second, the screen can only get smaller if the dot pitch and pixels(individual dots that can be turned on and off on a screen) physically get smaller in active matrix screens, or a new screen technology becomes widely available which has tighter dot pitch. The reason for this is, the C1VN screen is 1024x480. It is wider than 640x480, but 640x480 can be displayed in the center of the screen. If the screen was smaller(and therefore incapable of displaying 640x480), it would be impossible to install an OS on the machine, because a 640x480 display is the minimum requirement for all OS's. It just so happens on this machine that the screen is about the same size as the keyboard, so for practicality's sake, two out of three dimensions have reached their minimum size. The thickness of the screen couldn't be reduced anymore without putting the screen at greater risk of cracking.

Finally, the bottom half of any laptop can get no thinner than the thickness of the keyboard plus that of the thickest component. Upon inspection of the C1VN, the area under the keyboard holds space for the motherboard, 1.3" hard drive, one PCMCIA and Memory Stick slot, and all of the ports. The total unit is about 1" thick when closed.

Given all of this, it is highly unlikely that laptops will get any smaller and still remain comfortably usable in the not too distant future, unless we are somehow freed from the need for a laptop screen. Potential for this may come from improved versions of products like Sony's Glasstron products, which put a virtually huge screen into a currently costly pair of glasses that may be worn. The best resolution of the Sony Glasstron products currently only allows viewing 640x480. It would be pretty rough to try to get by with such a low resolution when using a computer, but it certainly could be done. These products are more directed at the viewing of DVD's and movies in a more personal and portable format than they are in providing a complete substitute for a computer screen or monitor of some kind.

As far a cheaper goes, one should expect to be able to buy something equivalent of the C1VN for $1000 within 2 years on the used market. However, I'd expect Sony's next version of the Picturebook to retail for the same or more money than today. Larger, less practical notebooks can definitely be had for less today, but generally will leave the owner scrambling for battery power in some way. The cost of batteries does not look like it is coming down anytime soon, as it is a well accepted fact that battery technology has fallen far behind the rest of technology. Lower power solutions like the Transmeta Crusoe CPU are where most advances are being made in terms of power supply and consumption.

1.1.1.1.12 What other features should I look for in a laptop to be used for recording?

A: 100BaseT Ethernet ports are always a plus on a laptop for getting music files transferred off for archival and trading. Firewire ports and devices are becoming very commonplace and are said to be eclipsing the value of USB ports. IEEE1394, a.k.a. Firewire or iLink400, is capable of 400Mbps of bandwidth, compared with 12Mbps of USB. This translates to approximately 50 megabytes per second of bandwidth on the Firewire port, enabling lightening-fast transfers to all kinds of peripherals like external Firewire hard drives. If your operating system supports it (currently only Windows ME), you can hook two laptops together and transfer files over Firewire at this same lightning speed, provided that your laptop hard drives can keep up with the available bandwidth of Firewire. Also, a relatively decent built-in soundcard and a headphone jack will allow you to casually listen to your recordings, albeit without all of the super fidelity your recordings will offer.

1.1.1.1.13 The screen looks small on the Sony Vaio C1VN. Shouldn't I look for a laptop with a larger viewable area?

A: With virtual screen sizing, you can set your display to 1280x1024. Virtual screen sizing works by automatically scrolling the entire screen depending on mouse position and chosen resolution. A good analogy to this is the use of a magnifying glass to read printed text. Quite realistically, 1024x768 or 1024x480(without virtual screen sizing) are totally sufficient for field recording. It's not hard to get used to virtual screen sizing, and it's still large enough to watch your favorite DVD at 640x480 if that's what you're into. While larger screens are certainly easier to read, they only come with larger, heavier machines, and they do stand an even greater chance of getting damaged with rigorous field use. Furthermore, it can generally be said that the larger the screen is, the more costly is it to repair.

1.1.1.2 Mac

1.1.1.2.1 What models of portable Macs should I consider?

A: The G3 and G4 are gaining a lot of attention lately as field usable options. Mac's implementation of the Firewire (IEEE-1394) is certainly making these machines quite desirable although software companies and hardware designers are not offering as much support for Mac products as they do for PC. Ironically, some of the world's most highly respected software, namely Pro Tools, is native to Mac. Obviously, these machines stand as viable options. (Mac users, please make submissions with your experiences so we can incorporate your experiences here.)

1.1.1.2.2 What are the minimum CPU requirements needed for accurate 24-bit recording?

A: (Mac users, please make submissions with your experiences so we can incorporate your experiences here.)

1.1.1.3 Sound Cards

1.1.1.3.1 What are my options for interface types?

A: There are currently two primary options for sound card interfaces: PC Card/Cardbus/PCMCIA and IEEE1394/Firewire/iLink400. For 24-bit recording, USB is not an option, except for the USBPre from Sound Devices (www.usbpre.com or www.sounddevices.com), which has 24-bit recording, but only 16-bit playback. See 1.1.1.3.10-14 and 1.1.1.4.8 for more on the USB interface.

1.1.1.3.2 What brands and models of soundcards are available for portable 24-bit field recording?

A: The Digigram VXPocket V2 ($600.00 street price), currently available, and the RME Hammerfall DSP with a Digiface or Multiface interface ($1000 for the card with a Digiface), said to ship May 2001. Both are PC Card/CardBus interfaces, but the RME Hammerfall DSP interfaces require external units that need separate power. Both have drivers for Windows 98/ME, Win2K, and MacOS.

1.1.1.3.3 What solutions are available to allow a half or full length PCI card to operate on a laptop for recording?

A: Magma (www.magma.com) makes a PCMCIA card with a cable connecting to an external one or two slot PCI card chassis weighing less than 2.5 pounds. This product would theoretically allow a user to put a multitrack PCI soundcard such as a SoundScape Mixtreme TDIF interface into a mobile unit and allow multi-track recording. No such tests with this product are known to date for that purpose, and it would be surprising to most if an internal laptop hard drive proved to reliably keep up with a full 8 tracks of 24-bit 48kHz recording. External Firewire hard drives may be the solution to this, but to date the author has found no Firewire hard drive solution capable of being easily powered by pure 12VDC. All solutions seem to require 120VAC (by built-in chassis supplies or external wall-warts), providing transformation and regulation to 12V for the mechanics and 5 volts for the circuitry. A complicated custom power supply would be required to make Firewire drives work in the field where power is uncertain.

A major benefit of the one-slot version is that it is capable of being DC-powered by a rechargeable 12Volt battery (not included) like the Eco-Charge Gamma series batteries. However, Magma does not recommend their product for laptops utilizing the Ricoh CardBus Controller, which unfortunately includes the Sony Vaio Picturebook C1VN.

1.1.1.3.4 If my laptop has two PCMCIA slots, can I use two Digigram VXPocket v2 cards and record 4-track instead of 2-track?

A: According to Digigram, because of a limitation in Windows ability to talk with more than one PCMCIA soundcard at a time, 4 tracks cannot be recorded utilizing two VXPocket cards. Definitely a bummer. Digigram may soon be releasing a 4-track version of their card, but it is unknown whether it will provide two distinct S/PDIF inputs, especially since use of two S/PDIF inputs at the same time would require that their clocks are in sync with each other, unless the S/PDIF ports are capable of resampling on the fly.

1.1.1.3.5 The Digigram VXPocket v2 card is pretty costly with a street price of nearly $600. Do I have to spend this kind of money to buy a 24-bit soundcard for my home machine as well?

A: Not necessarily. There are three options for working at home. If you like, you can purchase a PC Card reader for your desktop PC, and use the same VXPocket card at home. The DataChute PCI product($99, see http://www.antec-inc.com for details) has been used successfully to provide a desktop machine with 2 PCMCIA ports enabling use of a Digigram VXPocket card. Obviously this is a small hassle having to hook up the digi-in and digi-out on the VXPocket breakout cable to your home setup, but it might be worth the savings. This solution works great if you have a much more powerful workstation at home that you would prefer to do editing on. Option B is to use your laptop as a home machine. Simply plug in a keyboard, mouse, and a monitor, you can turn your laptop into a respectable home machine. Add an external Firewire hard drive and you're in the editing business! For those laptops with only USB ports, you can purchase a USB keyboard, a USB mouse, and a USB hub for a nominal cost. If you just want to play the 24-bit files residing on your home machine, and your laptop has room for an Ethernet card as well as the Digigram VXPocket, your third option is to use your laptop to map a network drive to your home machine, and play your files from your home machine on the laptop. Of course another option is to purchase a less costly PCI soundcard instead for your home.

1.1.1.3.6 My USB soundcard says it has 24-bit A/D converters, but is only capable of making 16-bit recordings. What's going on here?

A: For some reason or another, it appears that most USB soundcards seem limited to 16-bit recording, even though they utilize 24-bit A/D converters. One can speculate that 24 bits are captured by the A/D, and any gain applied is done in the digital domain. The 24-bit word is then either simply truncated to the 16 most significant bits, or some form of mysterious dither is applied. Needless to say, USB solutions with these specifications are useless to the 24-bit field recordist. Sound Devices, on the other hand, indicates that the USB Pre 1.5 can deliver 24-bits through the USB interface using Windows 98Se and 2000. See 1.1.1.4.8 for more info.

1.1.1.3.7 Can't I just use the built in sound card on the laptop?

A: Built-in laptop soundcards tend to have extremely poor fidelity and are much more subject to noise and interference from internal computer components. In short, built-in soundcards are about as good as computer speakers.

1.1.1.3.8 Can I record more than 2 tracks (multitrack recording)?

A: Yes! MH Labs is about to release the Mobile I/O 2882, which is an 8-channel 24-bit/96kHz IEEE1394 based(Firewire/iLink) device, complete with 8 microphone preamps, phantom power, and a 24/96kHz A/D converter with very respectable specifications (at least 110dB S/N ratio), in a portable form-factor about the size of a full-size laptop. This product will support AES/EBU in/out, Coaxial and optical S/PDIF in/out, and ADAT Lightpipe in/out, in addition to 8 analog-outs. It also supports Wordclock(1x) and SuperClock(256x) in/out, so you can sync it with other devices or additional Mobile I/O products. It remains to be seen which laptops can keep up with 8-track(or more) recording, as processor and internal hard drive speeds pale in comparison to that of the average desktop computer. The Mobile I/O promises full support for both PC and Mac, and includes both ASIO 2.0 and Wave drivers. Multiple Mobile I/O devices can be used concurrently for up to 128 tracks of 24/96kHz recording. The product is not currently shipping, but is currently priced at $1500(without onboard DSP's) and the manufacturer is taking $500 preorder deposits today. For folks requiring speedy editing capability, a version with onboard DSP's will be also be available for an additional $700 ($2200 total). While the specifications seem very good on paper, only time will tell as to how much of a true performer this product is. According to pictures on their website, the product requires 15W of 9VDC power, which may allow this unit to be run for 4+ hours on a field standard 12VDC 7AH battery(a 12V to 9V adapter may or may not be needed). See http://www.mhlabs.com/mobileio/mio.html for information on this breakthrough product. See 1.1.1.3.3 & 1.1.1.3.4 for more information related to multi-track field recording. Be on the lookout for new products from Digigram that may allow 4-track recording to a laptop.

1.1.1.3.9 Everything looks like it's set up properly. I've set my software to record from the digital in on the soundcard, but it says there's a problem when I hit record. What's wrong?

A: Most of the time when the operating system pops up a message the second you hit record or monitor, it is either because there is no digital signal or clock source present, or you have your clock source set up incorrectly in software.

1.1.1.3.10 Can I use a Roland UA-30 USB soundcard for 24-bit recording?

A: While we have had mixed feedback about this issue, a phone call to a seemingly knowledgeable tech support at Roland's subsidiary manufacturing company for the product, called Edirol (www.edirol.com) yielded the following information. They are aware of the widely misquoted specifications of the UA-30 on many websites that erroneously state that the UA-30 is in some way 24-bit device. Truth is, the A/D converters on the UA-30 are 20 bit, not 24 bit. Furthermore, because of a limitation in the USB port's available bandwidth in their chosen implementation, only 16 bits are written to the hard drive. There is no dither or noise shaping employed to reduce the 20 bits supplied by the A/D to the 16 bits that actually get recorded. The 20-bit words are simply truncated to 16 bits. We are seeking further details about this device but it appears to be inappropriate for 24-bit applications. The UA-30 therefore is incapable of supplying greater than 96dB of dynamic range (the maximum theoretical dynamic range of a 16-bit device), and in reality, because of the raised noise floor created by the distortion induced by truncation, the signal-to-noise ratio is significantly lower than 96dB. Additionally, as is quoted on their website, this unit draws 450mA of current, which is quite a hefty load for the laptop to supply through USB (maximum supplied by USB is 500mA), and would have a marked impact on overall power consumption.

1.1.1.3.11 Can I use the EgoSys Waveterminal U2A USB soundcard for 24-bit recording?

A: No. The U2A (http://egosys.net/eng/u24.html) has 24-bit A/D converters, but only records 16 bits to the hard drive. See 1.1.1.3.6, 1.1.1.3.13 and 1.1.1.4.8 for more info on the limitations of USB devices.

1.1.1.3.12 Can I use the Sound Devices USBPre 1.5 Microphone Pre-amp/converter for 24-bit recording?

A: Yes, see 1.1.1.3.13.

1.1.1.3.13 Can I use any USB soundcard for true 24-bit recording?

A: Yes. While USB v1.1 does technically afford the bandwidth (12Mbps) required for 24 bit/48kHz recording, due to bandwidth required for USB device polling in conjunction with and driver and OS limitations, true 24-bit recording is only possible in a couple of implementations available today. The best case scenario for an available USB audio device is the Sound Devices USBPre 1.5 (www.usbpre.com or www.sounddevices.com), which, according to tech support, is capable of sending 24-bit words from its own, or an external, 24-bit A/D through the USB port to the hard drive. However, on the PC, this is only possible with ASIO drivers, which bypass the 16-bit limitations imposed by the Windows Mixer where a separate software mixer is not supplied. Since the USBPre is actually only capable of 106dB of dynamic range, this equates to 18 actual bits used. This means that if you were to set your recording software to 24-bits, at least 6 of those bits would either be zeroes, or random digits. The USBPre does have only 20-bit output however, and truncates 24-bit playback signals without dither or noise shaping. Still, this appears to be the best choice in situations where portability and small form factor outweigh the need for the greater quality achievable by larger and more expensive solutions. It consolidates a microphone preamp with 48V of phantom power, a 24-bit A/D converter, and a soundcard into one, small, stealthy, rugged device, promising 12-15dB more than DAT can deliver for about the price of a small portable DAT deck.

Media Assistance, a 1999 German startup company, makes the USB One for Win98SE/2000 and Mac use, which boasts 24-bit Coaxial and Toslink S/PDIF input [only] when using ASIO drivers, and 20bit A/D input and output on line and mic level stereo jacks. Phantom power is questionable, and there is no digital output. Use of this product for 24-bit recording is limited to programs that can use ASIO drivers, such as Steinberg products like Cubase, and Nuendo. The box also looks rather large for a field recording situation. See http://www.media-assistance.com/English/Products/USB_Audio/usbOne/usbOne.html for more details on the USB One.

For the Mac, only Griffin Technology makes a product called the iMic. This product is a small (2 ounces) round case with two 1/8" stereo jacks (analog input and output), a mic/line level switch, and a 24-bit A/D converter. There is no digital input or output on this device. At only $35, one must seriously question just exactly how good this item can be. It looks downright scary. However, if you are a Mac user looking to do something with your laptop at a concert, this might sound as good as a small, portable consumer-grade deck. An audiophile would probably not want a copy of the show though. According to Griffin, the iMic "provides significantly superior audio input and output performance over built in audio. The iMic is a must have product for people who are serious about getting high quality audio in or out of their computers."

See http://www.griffintechnology.com/audio/imic_main.html for this information on this product.

1.1.1.3.14 Are there any caveats to the use of USB audio recording devices?

A: There are a few issues with USB as a recording interface. Power for these devices is usually drawn from the USB port, which is capable of supplying .5A. While often it is found convenient not to have an outboard device require separate power, the additional drain on a laptop battery may be of some concern, especially of the device provides phantom power for microphones. Second, proper operation of a USB device requires Win98 Second Edition, or Windows 2000. Although it has worked for some, Windows ME is not recommended. Win 95A/B and Win98 have known problems in handing USB ports and devices properly, and will only cause trouble. The use of Windows 2000 instead of Win98SE is highly recommended, especially with USB devices, because of better overall system interrupt handing. Additionally, USB ports usually get assigned a lower priority IRQ, which can make USB devices more prone to interruptions caused by other system devices while recording. Finally, a USB port being used for recording should have no other devices plugged in other than the audio device, because a non-audio device can steal USB bandwidth, power, and CPU cycles from the recording process.

1.1.1.4 General Laptop Questions

1.1.1.4.1 What are the benefits of using a laptop for field recording over any type of dedicated recording device like hard drive recorders?

A: You can visually see what you recorded, and you can scan for errors and glitches. You have almost instantaneous recording from the moment you hit the record button. You have a complete digital audio workstation in which can edit with any software you like, so you can master, fade, downsample, reduce word length (dither), track, archive, trade, and post your recordings in minutes while on the road. You have many more ways for offloading your files. You can transfer files via the internet back to your primary workstation. Your buddy can get a copy of the show on his laptop within 20 minutes of hitting the parking lot. Your favorite guitarist can have a copy of his previous night's performance on CD even though you are miles from home. You can trade shows in minutes while sitting around waiting for the show to begin, or at set break. And most importantly, you have a device you can use for something other than field recording.

1.1.1.4.2 How well does a laptop perform in inclement weather situations?

A: While no device likes getting wet, a laptop in some ways is less likely to have problems than a DAT deck. There is no dew light to come on, and no tape to stick to heads and guides. If it's raining at an outdoor show, keep the laptop in a clear plastic bag, preferably a large Ziplock. Put a good sized pouch of dessicant in the bag with the laptop if you have some. If you know it's going to rain at the location, plan on making extra room on your hard drive for your recording that day. Start recording early, close the lid if you can, and deal with it only if necessary. It's not like you're wasting tape. You can always delete the junk at the beginning of the recording with basic editing software later on after the show.

1.1.1.4.3 I've had trouble seeing the screen in bright or direct sunlight. What can I do about this?

A: First, keep the screen as squeaky clean as humanly possible. Sunlight will reflect off dust and oils from fingerprints and make it difficult to see the screen. I have used low water content isopropyl alcohol and a dustless cloth to clean my screen in the past. Your mileage may vary, do this at your own risk. Test a small area of the screen before applying any solvent full scale. Cranking up the brightness on the screen always helps, but can comp. Finally, during off hours, experiment with different color schemes and contrasts for your windows and background in direct sunlight to see what works best for you and your screen. And if none of those solutions are satisfactory, line the inside of a white pillowcase with thick black felt and put the laptop in the pillowcase. Do not under any circumstances run the laptop in a bag with dark colors on the outside in direct sunlight, or you risk cooking your laptop.

1.1.1.4.4 Where will my laptop reside in a patch chain if I decide not to run my own microphones?

A: With a digital distribution amplifier, or "patchbay," you should be at the front of the chain. If you don't have a patchbay, you're probably only going to be able to get on the end of a 16-bit DAT chain, and you will only get a 16-bit recording. If you want to take a 24-bit feed out of a Benchmark Sonic AD2K+ being used only for 16-bit recording, offer up a patch bay to help the 16-bit chain, take your patch directly out of the auxiliary digital out on the AD2K+, and ask the AD2K+ owner to set his auxiliary digital out to 24-bit. The need for a digital patchbay cannot be stressed enough if you are a laptop patcher, especially if you are looking to patch from a 24-bit source.

1.1.1.4.5 Is it OK to provide subsequent feeds (a.k.a. "patches") out of my laptop?

A: No, for several reasons. The most important is if anything goes wrong you will most assuredly be the target of negative vibes. More importantly for yourself is the fact that in order to allow this, you will have to either monitor the signal while you record (provided your soundcard has low latency direct monitoring support in hardware), or playback the signal while recording. Either one will add to the consumption of CPU cycles to your machine while recording, and can possibly end up stealing the CPU from being able to keep up with the recording process itself. A digital patchbay is a laptop recordist's best friend for this reason.

1.1.1.4.6 If I buy a laptop for recording today, won't it just become obsolete in a year?

A: A laptop bought today that is sufficient for 24-bit field recording will be sufficient for field recording tomorrow. Like all computer equipment, there will be improved models, but the value of achieving the best possible recordings immediately should outweigh the losses of not having the fastest, sleekest machine two years from now.

1.1.1.4.7 What PCMCIA or PC card Firewire adapter should I get to add a Firewire port to my laptop?

A: Several Firewire product manufacturers, such as MOTU, have noted problems with certain chipsets used in Firewire or IEEE1394 cards. NEC chips have been known to have problems, while TI and Lucent chipsets have been known to be very compatible and stable. Companies that make adapters employing TI and Lucent chipsets are Keyspan (uses TI), ADS Pyro (uses TI) , and Sonnet Tango (uses Lucent). Companies with products utilizing the NEC chipsets that have been known to have problems are Orange Micro and Western Digital.

1.1.1.4.8 Will USB ever be capable of supporting bandwidth required for 24-bit 96kHz recording?

A: The next revision of the USB specification, called USBv2, will have 40 times the bandwidth of USB v1.1. Expect to see devices and laptops with USBv2 ports in 4th quarter 2001. See www.usb-audio.com for information on USB and its use for audio recording.

1.1.2 Analog to Digital Conversion

1.1.2.1 Can I use my 16-bit or 20-bit A/D converter to make 24-bit recordings?

A: You can not achieve 24-bit quality without a 24-bit converter, but you can set your software up to record 24 bits and record a 20-bit source (e.g. from an Apogee AD1000) with some definite benefit over 16 bits. In this case, the least significant 4 bits will just be unused zeros, which amounts to some wasted space on your hard drive. However, there will be absolutely no gain over a standard DAT recording by recording a 16-bit source at 24-bits. You will only succeed in wasting 50% more disk space. This might only be a valid procedure if you intend on post-processing 16-bit audio in the 24-bit realm and you are trying to save yourself the additional step of 16 to 24-bit conversion.

1.1.2.2 If my 24-bit sound card has analog input, do I have to use an outboard A/D converter?

A: No, but chances are, as is the case with the Digigram VXPocket, an outboard A/D converter will provide up to 25dB more dynamic range than the analog in on the soundcard. Remember, 24-bit is now the industry standard for professional audio recording. Consequently, most manufacturers are developing products that utilize 24-bit A/D chips. However, achieving 24-bit performance requires an excellent analog stage and top-quality shielding from noise (a major concern inside of a computer environment). Not all "24-bit" devices sound the same, and buyers should be careful of manufacturers' use of the phrase “24-bit”, since a 24-bit system recording “self-noise” will only serve to record noise more accurately.

1.1.2.3 What are my current options for external portable (DC Powered) 24-bit A/D?

A: There is currently only one option for external portable 24-bit A/D: The Benchmark Sonic AD2K+. See http://www.benchmarkmedia.com or http://www.sonicsense.com/ for details on this well designed product.

1.1.2.4 Do I need to worry about truncation, dither, and noise shaping when recording 24-bits to the hard drive?

A: Dither and noise shaping is only required to mask digital artifacts when converting from a higher word length to a lower word-length (e.g. from 24 bits to 16 bits). When a 24-bit source is recording at 24-bits, there is no need for conversion or word-length reduction. Hence, there is no need for dither or noise shaping. Always make certain that the A/D converter is set properly, that the software and soundcard that you are using actually support 24-bits, and that they are set up for receiving the 24-bit feed.

1.1.2.5 I'm not yet entirely confident in the use of a laptop as my only method of recording. What are my options for simultaneously recording a 16-bit DAT for backup?

A: Amazingly, the Benchmark Sonic AD2K+ has the capability to provide a total of 4 digital outs in two channels: main and auxiliary. The main channel can be set to provide a dithered or noise shaped 16-bit signal, while the auxiliary simultaneously can pass 24 bits unadulterated. A DAT machine can be attached to the main channel, and the laptop can be attached to the 24-bit auxiliary output. If you are receiving a signal from a single 24-bit source, a distribution amplifier ("patchbay") in front of you can split the 24-bit signal, so your laptop can run parallel to a DAT. The only caveat here is that the DAT, if sent 24-bits, will truncate the last 8 bits and not add proper dither or noise shaping to smooth out quantization noise caused by the truncation. For this reason, the DAT deck should still be placed somewhere in the dithered 16-bit chain.

1.1.3 Hard Drive Recorders

1.1.3.1 Who makes hard drive recorders that support 24-bit word length?

A: Both Tascam and Mackie make 24-bit 24 channel hard drive recorders. These units are rack mount, and provide up to 1 hour of continuous 24 channel recording. Zaxcom and Nagra also make 24-bit hard drive recorders. And there are a variety of short-duration hard drive recorders out there aimed at musicians and composers for the purposes of songwriting and creating demos, although most of these products are not concerned with the quality recording 24-bit offers, despite frequently boasting 24-bit A/D conversion, and are really only on par in sonic performance with 16-bit devices.

1.1.3.2 Are there any portable hard drive recording solutions?

A: There are currently no reasonably priced ($2500 or less) 24-bit recording solutions available to the field recordist. The Zaxcom Deva II is a 4 channel, 24-bit capable hard drive recorder with a built in mixer, time code, and onboard 20-bit A/D and D/A converters, priced at $9950. The website http://www.zaxcom.com/audio/deva.shtml states 4.0 track hours per disk, and contains one internal disk. External SCSI disks can also be attached for additional track time. It weighs about 5.1 lbs with a 2 hour battery. This unit is highly regarded in the film industry as a standard, but its extensive feature set along with its price make it a bit foreboding to the recordist on a budget.

1.1.3.3 How long can I record continuously with a hard drive recorder?

A: That depends on the size of the hard drive, the number of tracks recorded, the word length, and the sampling rate. The formula for calculating this is as follows:

(Hard drive size in GB * 10E+9)/((word length in bits * # of tracks * sampling rate in Hz * 60)/8) = record time in minutes.

For example, for a 9.1 GB drive, recording 2 tracks at 24 bit, 48KHz, you have

9,100,000,000 / ((24 * 2 * 48000 * 60) / 8) = ~526.6 minutes or 8.78 hours

This same formula can be used to calculate available record time on a laptop as well.

Again, note that a 9.1GB drive is really only 8.68GB. This is because drive manufacturers call a 1GB=1000MB, but in reality, 1GB is 1024MB, 1MB is 1024KB, and 1KB is 1024 bytes. So are they misrepresenting their products? Yes. Are they getting away with it? Unfortunately, the answer is yes as well.

1.1.3.4 Why haven't we seen many portable hard drive recorders suitable for 24-bit field recording yet?

A: Besides the obvious fact that technology has only recently made it possible to create a viable hard drive recorder, there are several factors involved. Firstly, the majority of the need for 24-bit recording has classically been in the studio, not the field. Additionally, many of the smaller studios are only now beginning to convert their digital audio work to the 24-bit realm. There are still a large number of audio engineers who prefer to deal with the analog world because it is still a great sounding medium. Other lack the understanding, and in some cases even fear digital audio and computers required to deal with it. Furthermore, working with 24-bit digital audio requires studios to purchase new 24-bit capable workstations and large 24-bit digital mixing consoles, which are far more costly than even the best classic analog mixing consoles. Basically, there is no point in recording 24-bit audio in the field if you can't deal with it in the studio. Therefore, the market is currently focused on 24-bit products for the studio, with little attention given to the field.

There are two other major factors in all of this. Because of the rapid rate of increase in capacity of hard drives, and decrease in cost, manufacturers have been unwilling to base an expensive product on a given hard drive capacity that will certainly be obsolete very shortly. If they put a biggest drive available in their unit, it will be very costly, and even more costly to be able to swap drives out with other drives. Therefore, lesser capacity drives are typically used, like 9GB drives (even though much larger drives exist) because they make up a cheaper solution for studios purchasing multiple drives.

Secondly, the majority of digital recording in the studio is usually targeted at short durations of multiple tracks, not long durations of 2-track recording. Unlike the field recordist, the studio has little need for 4 continuous hours of 2-track recording, and again, equipment manufacturers have been focusing their products largely to the needs of the studio. Hopefully, all of this will change, but it remains to be seen.

1.1.3.5 What are the obstacles manufacturers of hard drive recorders face in bringing products to market?

A: Recouping costs before obsolescence - One major obstacle in bringing an affordable portable hard drive recorder to market is the need to avoid obsolescence before sales of the product hits its make or break point. Like all products, development, production costs, and marketing and advertising costs must be quickly recouped in order for a product to be successful. This is a problem in a machine designed around the capabilities of a hard drive, where hard drives double in capacity every 18 months, and prices fall on a daily basis. Additionally, for the sake of portability and reduction in power consumption, the smaller and significantly more costly 2.5" hard drives popular in laptops are better suited for hard drive recorders than the 3.5" found in most computers today. The smaller drives are also much more shock resistant. A portable hard drive recorder manufacturer has to base their product on a production run of initially relatively expensive hard drive, and adds to the retail cost of the unit. It may take a manufacturer the sale of 3000+ units before they break even on the product. Because of a relatively small market for such a product, compared to the market available to products like MiniDisc and MP3 recorders, a manufacturer's efforts to develop a product that is likely to be rapidly obsolete would be a great risk. Some might speculate that such a business risk could only be taken by large capital companies like Sony, which may also already have somewhat competing products.

Interface types - Another obstacle is the choice and implementation of an interface used to offload hours recorded music. Most comparable products today are aimed at the MP3 market, and exist for every good reasons. An average MP3 song is 4-7MB in size, and most devices available today have USB ports on them used to transfer music to and from the device when attached to a computer. While USB technically offers 12Mb maximum bandwidth, reality is more like 500KBytes per second maximum transfer rate in actual use. Computers with other USB devices attached may not even have that much available for music transfers. Consequently, an average song can be transferred to a portable MP3 player in about 10-15 seconds. However, 3 hours of 24bit 48KHz music (a single recorded concert) is about 3GB, almost 1000 times the size of an MP3 song, and at that rate, it would take 30-40 minutes to transfer one show to a computer with USB. Imagine how long it would take to transfer a day's worth of festival shows at this rate? Therefore, more likely to be employed by hard drive recorders suiting the needs of the 24-bit field recordist are interfaces like IEEE1394 (aka Firewire or iLink400, depending on the implementation). IEEE1394 interfaces are definitely more costly to implement than USB interfaces, although it is certainly not prohibitive. Devices utilizing the IEEE1394 specification have generally been rather proprietary in their implementation, and require software drivers to be written for the specific computers and OS platforms, further adding to the development costs. While a few audio equipment manufacturers like Tascam now make both audio and computer solutions (like the Tascam USB428 soundcard), they do not have huge teams of software and driver developers at their disposal.

100Mb Ethernet may seem like a possible solution as well, but this implies a network stack like TCP/IP is employed, requiring a method of IP configuration for the unit. Not everybody is network savvy, and most people do not have a DHCP server in their home. The use of Ethernet definitely complicates matters, and its implementation in a hard drive recorder would likely approach the implementation of what is known as embedded computer systems. Embedded systems implementations are relatively new, and very costly. Basically, a device with an embedded system contains the equivalent of a full computer stripped down, without the standard screen, keyboard, and ports we are accustomed to. Start putting more computer components into a hard drive recorder, and you might as well just use a laptop. Manufacturers are aware of this.

Stand alone or computer interfaced? - What would certainly be nice is a portable hard drive recorder with cheap, removable/swappable hard drives that aren't so quick to become obsolete. The question then becomes, what will we do with the data on these drives? Where will we edit and process them? Will they be stored on the hard drives permanently, or are they destined for some other archival medium. If editing capabilities are added to a hard drive recorder, surely the cost of the unit will skyrocket. How do I make a copy of the material for someone else, or for myself as backup? Again, as these questions are answered, the need for the involvement of the computer seems clearly evident. Contrast this to the needs of the average 16-bit field recordist today. A DAT or MiniDisc recorder is a stand-alone solution, and users may never need a computer involved.

Power - Here's some general facts about hard drives: The smaller the hard drive, the more it costs. The smaller the drive, the less capacity it is able to achieve. Finally, the smaller the drive, the less power it requires.

The market for potential manufacturers of 24-bit hard drive recorders are well aware of the fact that every day new soundcards in a portable factor are made available for an equally increasing number of small and viable laptops. They may not find it profitable to produce such a product in an environment with rapidly growing competition.

As you can see, there are many logistical, technological, and economical obstacles today, impeding progress toward the development and release of the holy grail of 24-bit hard drive recording. Again, hopefully this will change in the future, but given all of this, it sure doesn't look like a <$2000 stand-alone portable hard drive recording solution is on the horizon anytime soon.

1.1.3.6 What are the obstacles hard drive recording users face related to their use in field recording?

A: The biggest problem a standalone portable hard drive recording user will face in the field is how to cheaply and conveniently archive the recordings, especially while on location for extended periods. The same obstacles are faced by laptop recordists. However, with a laptop, you have many more options available to you for offloading recordings, including CD burning, Firewire and PCMCIA hard drive solutions, and high bandwidth connections to the Internet.

1.1.3.7 It seems like both Tascam and Mackie have had multi-channel (24 channels of 24 bits) hard drive recorders available for a while. Why shouldn't I expect them to become available in a portable format?

A: Because these units are designed to be used in a studio environment. Any need for a multi-channel hard drive recorder with editing capability will also be commensurate for a need to power the many devices supplying signals to the multi-channel recorder. In other words, the availability of A/C power can be easily assumed with these recorders.

1.1.4 High Resolution Digital Audio Tape

1.1.4.1 What is High Resolution Digital Audio Tape?

A: High resolution DAT tape is a proprietary 24-bit tape format invented by Tascam and used in the Tascam DA45HR. The "HR" stands for high resolution. 24 bits are stored on a DAT tape by running the tape twice as fast as 16-bit DAT tape. These 24-bit DAT's are not backward compatible with 16-bit DAT decks, although the DA45HR retains the ability to play and record 16-bit DAT.

1.1.4.2 How reliable is high resolution DAT tape?

A: Good question. Probably about as reliable as 16-bit DAT tape.

1.1.4.3 Should I expect to see a portable high resolution DAT recorder in the near future?

A: According to verifiable sources, the production of DAT transports is being being cut off, and equipment manufacturers like Tascam say they have enough transports left for the next 2-3 years. After that, they are no more. Since Tascam is the only affordable manufacturer of a 24-bit DAT machine, and its format is proprietary, not many have been sold. Therefore, it would appear unlikely that Tascam will ever be coming out with a portable version of its weak-selling deck.

1.1.5 High Resolution DTRS Tape

1.1.5.1 What is High Resolution DTRS Tape?

A: DTRS stands for Digital Tape Recording System. The media used is 8mm magnetic tape very similar to Hi8 video tape, but differing in the improved tape formulation used. It is a multi-track recording solution, usually 8 tracks of 24-bit 48kHz audio. Some implementations allow 24bit 96kHz recordings to be done on 4 tracks instead of 8. Tascam and Sony both make high-resolution DTRS tape solutions, although the technology was invented by Tascam. DTRS is aimed at the same market ADAT dominated for years, but DTRS is superior to ADAT in reliability, because the 8mm tapes are less prone to stretch and wear than ADAT tapes (VHS format). DTRS machines were first introduced by Tascam in 16-bit format in rack-mount machines like the Tascam DA38, DA88, and DA98. Tascam later followed up with a 24-bit version of their machine, called the Tascam DA78HR. The "HR" stands for high resolution, synonymous with 24-bit recording. With the DA78HR, Tascam was able to achieve 24-bit recording on the same tape as 16 bits without sacrificing record time, by increasing bit density on the tape. The Tascam DA78HR can be seen in the racks of sound engineers of popular bands like The Dave Matthews Band and String Cheese Incident, who multi-track all their performances. Multiple DTRS machines can be combined and "slaved" off each other for simultaneous synchronized tracking of 40 tracks or more.

1.1.5.2 How reliable is High Resolution DTRS tape?

A: The Tascam manual for the DA78HR recommends the highest grade of metal formulation 8mm tape be used for 24-bit recording. Tascam does make note several times that 24-bit recording is more prone to error than 16-bit recording. They stress the importance of clean tape heads, stating that just because heads are clean enough to yield an error-free 16-bit recording, a 24-bit recording can suffer more easily from dirty heads and transports. Fresh tape is always highly recommended for 24-bit recording. Proper and regular maintenance of the transport is crucial to the reliability of 24-bit recording, because recording 24 bits on the same tape format as the older 16-bit machines requires operation at tighter tolerances. Still, a properly maintained 24-bit DTRS machine can definitely yield many reliable hours of operation.

1.1.5.3 Should I expect to see a portable high resolution DTRS recorder in the near future?

A: DTRS recorders are multi-track devices targeted to studio use. It is doubtful a portable version is very marketable, due to the lack of requirement of a multi-track solution where A/C power is not available.

1.1.6 High Resolution Compact Disc

1.1.6.1 What is High Resolution Compact Disc?

A: Alesis makes a stand-alone product called the Masterlink CD recorder for about $2000. It enables a proprietary implementation of 24-bit recording to standard CD-R's in a data format. On a 74 minute disc, 45 minutes of 24bit 44.1KHz audio can be recorded. The product is aimed at studio use for quick one-off demos of recordings while retaining the sonic quality of 24-bit recording.

1.1.6.2 Will high resolution Compact Discs play on any devices I may currently own, such as CD or DVD players?

A: Because the format is proprietary, these discs will not play in any devices other than another Alesis Masterlink CD recorder.

1.1.6.3 How reliable is high resolution Compact Disc?

A: Good question, probably as reliable as an audio CD-R, or perhaps even more reliable because it uses a data format (with more robust error correction) rather than an audio format.

1.1.6.4 How likely is it that we will see high resolution Compact Disc recorders become available in a portable format?

A: Not likely at all. Alesis is the only manufacturer, and the current product is sufficient for the target market.

1.1.6.5 Should I just wait for on-the-fly portable high-resolution CD recorders to become available?

A: No. See 1.1.6.4. Anyway, recording CD's of any kind on-the-fly is never a good idea. Dust, vibration caused by loud music and movement of the unit, and bad disks can all cause a CD burn to fail, leaving the recordist with nothing but a shiny coaster and tears running down their cheeks. CD burning is much better left to controlled environments that the field cannot offer.

1.2 Software

1.2.1 What software packages are available for 24-bit recording?

A: For PC's, Sonic Foundry makes Sound Forge 5.0 and Vegas Audio 2.0. Vegas Video 2.0 is the same as Vegas Audio 2.0, but much more expensive. See www.sonicfoundry.com for more on the Sound Forge 5.0 and Vegas products. Sekd (www.sekd.com) makes Samplitude 2496. Steinberg supports 24-bit recording in their products Cubase 32/VST 5.0, Nuendo, and WaveLab. See http://www.steinberg.net/products/index.phtml for more Steinberg products.

For Mac's, good question.

1.2.2 Why would I want to use multi-track editing software to record only two tracks?

A: Some multitrack packages may provide advantages in better disk management and lower CPU utilization over other two track packages. They are typically written with more optimization in mind, because they must enable as many tracks as possible to be recorded without dropping samples.

1.2.3 Before I trust my laptop for recording, what types of tests can I perform to ensure an accurate, lossless recording?

A: An excellent question, as proper testing is very important to the success of the laptop recordist. Recording a piece of material for 3 hours and listening back is certainly one way to test. However, this requires quite a bit of extended, intense listening, and it may be easy to miss an anomaly in your recording if you're not paying 100% attention. Also, it may not be possible to hear the absence of one or two samples per second upon playback. There is really only one way to test your setup with complete accuracy and assurance that you need to record with confidence. Use the following procedure:

1. To start, you must have some 24-bit material on another computer with a 24-bit soundcard. At the very beginning of the recording, place a “landmark”spike of some kind. Perhaps a one-sample “pop” at -10dB will do. You can create this spike with a tone generation or FM synthesis process in your editing software. Be careful when playing back this spike as it can damage your speakers if played too loudly. Place another of these “landmark” spikes at the very end of the recording. Call this recording “file #1”

2. You must be sure that the playback computer can truly play the 24-bit material with 100% accuracy. Playback the 24-bit material on one computer, and record it with the laptop, using a short, high-grade digital cable in as interference-free an environment as possible. Call this recording “file #2”

3. When the recording process is done, trim the beginning and ending of the recording so that the landmark spikes are once again the boundaries of the recording. Save the recording (file #2). In theory, the files should both contain the same exact number of samples, if everything is working properly. However, this is not enough of a test, because while the same number of bits may have been recorded, there may have been errors in the transmission.

4. Using a network or some method using swappable drive, copy file #2 from the laptop back onto the computer used for playback.

5. Now load file #1 and file #2 into the playback computer’s editing software.

6. Perform a 180 degree phase inversion process on file #2. This turns all the valleys into peaks, and all of the peaks into valleys.

7. Now mix file #1 and file #2 together at equal loudness.

8. If you achieved a bit for bit recording, the resultant wave file should be completely silent (the two waves should cancel each other out perfectly). All samples should equal zero. If not, you have a problem to resolve on the laptop.

9. Try this recording test while simultaneously moving the mouse around the screen, plugging in the power adapter, and especially while closing and opening the screen lid on the laptop. Stress test the laptop while recording to see how much non-recording related stuff you can do to it without dropping samples or stalling out. Learn what you can and cannot do while recording.

10. If one recording package performs undesirably, try a different package.

1.3 Techniques and considerations

1.3.1 The Analog Stage

1.3.1.1 How important are the analog stages to a 24-bit recording?

A: Extremely important. Constant, low level noise from interference or distortion, (inherent or otherwise) of any kind can mask the low level sounds you are trying to capture. The analog stage, while often overlooked, is as equally critical as the quality of the A/D conversion, if not more important. Without top-quality power supplies and microphone pre-amplification, the benefits of 24-bit resolution are severely compromised. While you may see many devices touting 24-bit analog-to-digital conversion, if the dynamic range is less than 144db, than they can not truly take advantage of full 24-bit dynamic range. Many devices boasting 24-bit (the industry's favorite buzz-word) have dynamic range specs. under 96db. It must be understood that analog circuitry is the weak link in the digital recording chain. The best analog equipment available with today's techology can only deliver a maximum of 120db dynamic range. An analog stage rated at 115db or above in the dynamic range specification is considered excellent and takes best advantage of 24-bit resolution.

1.3.1.2 What will happen to my recording if I use unbalanced analog connections?

A: Unbalanced analog connections are very susceptible to the induction of noise from the environment as well as the creation of ground voltage potential. Specifications of 24-bit A/D products like the Benchmark Sonic AD2K+ are all based on balanced analog connections, and to run a signal with an unbalanced stage in it will likely negate the benefits of 24-bit recording by lowering the signal-to-noise ratio of the total analog stage. Additionally, devices with unbalanced outputs, that can't provide levels above +12dBu, typically do not offer great enough dynamic range to justify 24-bit a/d conversion.

1.3.1.3 What kinds of situations can mess up 24-bit recordings?

A: The proximity of devices emitting low-level noise, such as cell phones, AC to DC power inverters, and CRT monitors can negate or mask much or all of the benefit of 24-bit recording over 16 bits. Even 16-bit recording is at risk of these induced interferences, although the noise will not be captured with the clarity 24-bit recordings offer. Flicking lighters in a concert taper section (my pet peeve) have a way of appearing with unsurpassed clarity on 24-bit recordings. Bring an extra penlight for your Bic-flicking flashlight-less friends, and tell them to get lost if they expect to use their cell phone in the tapers section. The Nextel speakerphone type cell phones are especially bad news for all tapers, as large amounts of noise are emitted from them even before they start ringing.

1.3.2 The Digital Stage

1.3.2.1 If I intend to make 16-bit 44.1kHz CD's from my 24-bit recordings, should I record at 44.1kHz or 48kHz sampling rate?

A: This depends on how you choose to think of the benefits of 24-bit recording. A 24-bit recording allows the recordist to capture more resolution than is needed for CD, and in turn allows for greater headroom and precision when processing audio. If you are looking for an easy, one step process (word length reduction with dither/noise shaping) to get your recording onto a CD, and this is your primary goal, then you may want to consider recording at 44.1kHz. However, two other facts should be seriously considered here. First, if you have 24 bits to work with, running a process like sample rate conversion required to convert from 48kHz to 44.1kHz ends up retaining far more precision than the classic 48kHz to 44.1kHz sample rate conversion performed in the 16-bit realm. This added precision reduces quantization noise typically associated with this process, and therefore makes the process far less damaging to the audio than 16-bit folks are used to. This is why you will see compact discs boasting the use of higher bit lengths in the recording process. Second, you'd have to ask yourself the following question: "Is my goal to have the best recording I can make for the future, the best recording I can make for today's common format, or both?" The future will bring playback mediums supporting 24 bits at 96kHz. In a sample rate conversion process, converting from 48kHz to 96kHz (also known as upsampling) requires little real math, as every sample is simply doubled. The process runs very quickly compared to a conversion from 48kHz to 44.1kHz, or even 44.1kHz to 96kHz. DVD-Audio specifications do state that it will be capable of 88.2kHz as well as 96kHz (and 44.1Khz doubles to 88.2kHz). Since 96kHz is bound to be the new, more popular standard (44.1kHz "CD quality" is unfortunately considered the common standard today, not 48kHz), you may want to perform your recordings in a manner that better lends itself to easy conversion to the best of common standards of tomorrow. To record at 44.1kHz from the beginning can arguably be considered "crippling" the capabilities of your recording for the future. So it is the lesser impact of processing, coupled with the desire to be more ready for the maximum capabilities of future media formats, that has the author recording everything at 48kHz, and downsampling to 44.1kHz in cases where the creation of CD's is necessary.

While the importance of CD mastering still remains high for production oriented-professionals, anyone who has the luxury of building an audio archive for the future would be well advised to keep sampling rates as high as possible and only make the compromise to 44.1kHz CD's when necessary for listening convenience.

1.3.2.2 What happens if I set my software to record at 24 bits, but my source is actually only 16 bits?

A: The last eight bits will be unchanging zeros, and you will use 50% more space. This would only be desired if the recordist plans to process 16-bit audio in the 24-bit realm, looks to save the additional step of 16 to 24-bit word length conversion after recording. The recording will sound like a 16-bit recording.

1.3.2.3 What is 32-bit floating point and how will it affect my recording and editing?

A 32-bit floating point is not of value for simply recording because of limitations in the dynamic range of analog stages like that in A/D converters and microphone preamps. However it can be extremely valuable when editing 24-bit material. As many functions such as equalization, crossfades, and normalizing require additional bits (increased word-length), a higher than the preferred word-length must be applied to prevent truncation between processes. Once the editing is complete, the file can be returned to its 24-bit format with perceptively little loss. **Watch for expansion of this section soon.

1.3.3 Levels

1.3.3.1 How important is it to get "as close to 0 as possible without going over" with 24-bit recording?

A: With 24-bit recording, that game is over! Headroom is among the greatest benefits achieved with 24-bit recording. Now you can set your levels low, capture wide dynamic range with better accuracy, and best of all, never clip or compromise detail. Every bit offers about 6dB of dynamic range. A recording with peaks hitting -12dB maximum will still take advantage of 22 bits worth of dynamic range and resolution. The same low level recording in the 16-bit world would yield a 14-bit recording (yuck!). The 24 bits of precision again allows the recordist to normalize such a recording to higher listening levels with minimal impact to the quality of the recording. This is one of the greatest benefits 24-bit recording offers. Greater quality of lower level signals, without any clipping. This age-old 16-bit level game will probably be the greatest habit to break for the new 24-bit recordist. Where there are no 16-bit devices down line from your 24-bit recorder, set conservatively, and forget it. Even at –48dbfs, you are still capturing resolution as great as a 16-bit recorder at 0dBFS.

1.3.3.2 How can I set my levels to maximize the value of 24-bit recording while continuing to provide 16-bit devices with a good signal?

A: Most 16-bit recordists have a tendency to set levels either too conservatively, or not conservatively enough. Many times, peaks throughout one section of a recording hit -1dB, and then soundman turns it up a little for the second set of a concert, and the over-confident, overly comfortable recordist brickwalls and clips peaks repeatedly in the second set. Efforts to avoid or fix this situation may have the recordist constantly adjusting levels (called "riding the levels"), which then has the effect of limiting or compressing the dynamic range of the recording. 24-bit recording allows the recordist to leave some headroom for increasing levels without causing the recording to suffer in quality. When recording with such headroom, even associated 16-bit decks are too terribly compromised, provided that the analog-to-digital converter can send a 16-bit feed simultaneously with proper word length reduction with dither and noise shaping. This is a feature of the Benchmark AD2K+. Setting levels where maximum peaks are generally at -6dB is a safe bet, leaving headroom while not making it necessary to normalize levels in later production. When the soundman cranks it up 4-5dB in the second set, and you're out in the front-of-board crowd shakin' your bones because you are confident in the headroom you've allowed, your 16-bit patchers will thank you for a "perfectly leveled" recording, with no clipping, and no dynamic range compression. If the levels never increase after all, you can be assured that you got a great recording, and the 16-bit recordings didn't suffer too terribly. When recording sound effects or speech, you can enjoy the quite moments before the gunfire, and savor those whispers as much as the screams that your talent delivers.

1.3.3.3 Do I need to set levels my levels on the laptop?

A: No. There should be only one gain stage when recording, and that should be at the preamp. Recording though the digital input on the soundcard supports a single gain-stage scenario, and no digital gain enabled by the soundcard should be employed. Only when connecting microphones to the soundcard's analog input where no external preamp is involved (I don't know why anyone would want to do that, since this probably does not take advantage of the sonic improvement 24-bit recording has to offer) should some gain be applied.

1.3.3.4 Should I use a limiter in the recording process?

A: Absolutely not. Limiters are a form of compression which decrease dynamic range. To use a limiter is to defeat the very purpose of 24-bit recording, which is to provide greater dynamic range.

1.3.4 Sampling Rate

1.3.4.1 How practical is 96kHz vs. 48kHz currently?

A: An average 24 bit/96kHz concert recording requires 6.5-7GB of space, compared to 3-3.5GB for a 48kHz recording. 96kHz recording effectively doubles the time it takes to edit, archive, copy, and otherwise manipulate the audio, in addition to doubling storage costs while simultaneously cutting available record time in half. As hard drives continue to get bigger, faster, and cheaper, and higher processor and system bus speeds become available, 96kHz recording becomes more feasible. Today, many recording engineers still scoff at the immense space required to support 96kHz, not to mention 192kHz, and argue that the differences would be negligible at best. Still, once storage solutions and editing workstation performance improves (as we can surely expect soon), we'll undoubtedly see more recordists taking advantage of yet another possibility for sonic improvement.

1.3.4.2 My mentality has always been to record at 44.1kHz because I don't want to perform a sample rate conversion for CD's. Should I continue to record at 44.1kHz, or am I better off just using 48kHz for the future?

A: See 1.3.2.1.

1.3.5 Managing your Power Requirements

1.3.5.1 How can I power all of this equipment in the field?

A: In a word, batteries, batteries, and of course, batteries…and artful power management definitely helps. There is no substitute for DC power when it comes to improving sonic quality. Having AC power anywhere near your equipment can have a detrimental affect on your recording. This is only more true for 24-bit recording, where lower level signals like AC hum can be reproduced with greater clarity. For laptop recording, sufficient amount of power from internal laptop batteries is most desired. Still, if AC power is available, you can take advantage of it, but always have the internal battery in the laptop ready to go in case of a power failure. Again, keep the AC power source as far away from your other recording equipment as possible. The self-reliant field recordist does not count on AC power availability. For a creative alternative power solution and a better understanding of laptop power requirements, see section 1.3.5.3.

1.3.5.2 Can't I use a 12V power inverter to supply my devices with 120V?

A: 12 Volt power inverters all create considerable noise from induced electromagnetic radiation (from a process called "waveform chopping"), and this noise and distortion can easily mask the benefits of the low-level resolution 24-bit recording offers. The noise can get into your microphone and line level cables and all analog circuitry, no matter how well shielded you believe your system to be, leaving you with a great representation of noise on top of your intended recording. In some cases, your recording can easily be much worse than if you had just stuck to 16-bit DAT recording from a pure DC power source. Additionally, while DC>AC inversion sounds like a good idea, you are still dependent on the DC source so you might as well utilize it in its purest form whenever possible. It should also be noted that there is significant power loss in any conversion from 12DC -> 120VAC -> anything DC.

1.3.5.3 Is there any way I can use lead-acid batteries or gel cells to power my laptop?

A: Early 24-bit field recording to a power-hungry laptop was successfully enabled by modifying the connector(cigarette lighter adapter) on a 12V laptop auto adapter to connect to 4-pin XLR connectors on standard 12volt 7AH lead-acid batteries commonly in use by recordists today. This adapter is a DC-DC converter, and does not have the noise and interference problems associated with low voltage DC to high voltage AC conversion. Using a 12V auto adapter, the laptop gets its power externally, as if it were plugged into an AC outlet. Therefore, when using this rig, you will not get feedback on battery performance from the laptop. Therefore, the following methodology can be used to get an idea of how long you can run from an external 12Volt battery. Every laptop will have power requirements shown on a label on the bottom of the laptop. It will tell you both voltage and current(amp hours) requirements. Because laptops are DC devices, simply multiplying volts * amps will give you the power rating of the laptop in watts. Similarly, a new 12Volt 7AH battery is capable of supplying a maximum of 84 watts, assuming you beat the battery to death, which is never advisable for lead acid batteries. Running a battery down to 15-20% capacity is probably as far as you can go without permanently impairing the life and chargeability of your battery. So for a 12V 7AH lead acid battery, 65-70 watts is a safer estimate of usuable battery life. Additionally, the laptop auto adapter supplies voltages higher than the 12V input voltage by use of a transformer. This is not a particularly electrically efficient process, and you can probably estimate about 15-20% loss from this process due to heat, so 55-60 watts is probably even more realistic.

As an example, if your laptop says it requires 16 volts and 2.5AH, then you say the laptop requires at most 40 watts. We say "at most" because this number is about 5-10% greater than what the laptop would draw if everything on the laptop was at full use all at once. If your laptop has a CD-ROM or DVD player in use at full speed, a dead internal battery being charged, all PCMCIA slots and other ports like USB are in use, the screen is at full brightness, and the CPU is at full speed (for Transmeta Crusoe or Intel SpeedStep microprocessors), you can assume that the laptop requires no more than 40 watts to run at that moment. If you aren't using all of your peripherals (especially the CD ROM and screen) while recording, your power requirements can be somewhat less than 40 watts, perhaps more like 25-30 watts. In this case, given a power supply capable of 60 watts, you should be able to estimate longevity at 2 hours. Coupled with an internal battery lasting 2-3 hours, such a rig should be sufficient for a 3-hour concert recording. Of course, you may be able to run this rig for a little longer if necessary in a pinch, but it will come at the expense of decreasing the overall life of your lead-acid battery. As you get to the end of your external battery life, be sure you have your internal battery in your machine, in case the external battery all of the sudden cannot further power the laptop due to low voltage.

It should be noted that the author has never seen a laptop willing to run directly off of 12Volts. Laptops typically require DC voltage supplies between 15-18VDC. This is probably because voltages significantly higher than 12Volts are required for the laptop to assure the adequate supply of the nominal 12V for drive motors, fans, and the screen as well as the 5V for electronics.

1.3.5.4 While recording with a laptop, what can I safely do to conserve power without jeopardizing the recording quality?

A: If your CD-ROM or DVD-ROM unit is removable, remove it before recording. If you have a PC card in your machine other than the soundcard, remove it as well. Your laptop screen should have some form of brightness control. Lowering the brightness of the screen helps save a lot of power as well. Closing the screen, turning it off, is even better. Tests should be performed to determine if the act of closing the screen causes the machine to hiccup in any way. This was a problem on an IBM Thinkpad 600E, where closing the screen caused the power management processes to temporarily divert enough processing power from the recording process, resulting in a short but very noticeable period of dropped samples. In investigating this problem further, it was shown that some programs stopped recording completely when the lid was closed. This happened to both Sekd's Samplitude 2496 and Steinberg's Nuendo products, but not to Sonic Foundry Vegas Audio 2.0 or Sound Forge 5.0. In short, test, test, and test.

1.3.5.5 Which of the power management features should I enable on my laptop?

A: The best answer for reliable, uninterrupted recording is "none." This may not be the answer everyone is looking for, but it is nonetheless the case. The best practice is to disable everything non-essential manually, and turn off power management. The only exception to this may be in the case of a machine with the Transmeta Crusoe CPU like the Sony Vaio. In this case, after fully testing, you may find that enabling the dynamic CPU throttling feature safely provides some level of power conservation. Watch the system CPU meter and the Crusoe CPU monitor while recording. If you set your Crusoe CPU to "LongRun Power Management (Performance Mode)," and the Crusoe CPU monitor stays pegged for 90% or more most of the time while recording, it may be safer overall to just set your machine to "Maximum CPU Performance" instead. Note that even if you completely disable power management, a laptop will turn off its screen if you close the lid. Depending on the software and settings used, CPU performance may or may not be a significant factor. Sonic Foundry Vegas Audio 2.0d has been found to work reliably with CPU power management enabled, as the software manages its CPU utilization well.

1.3.5.6 What about the Galaxy Audio Faroutlet as a source of portable power?

A: The following is taken from the Galaxy Audio website. "The FAROUTLET Model 250S is a fully self-contained portable source of 110 volt, 60 Hz household current. It also serves as an uninterruptable power supply (UPS) for computers and as an automatic charger for 12 volt batteries (multiple batteries can be charged simultaneously). About the size of a lunchbox and lightweight enough to carry in one hand, the FAR OUTLET Model 250S provides up to 250 watts continuous and 400 watts peak power and operates for hours between charges (accomplished by plugging the FAR OUTLET Model 250S into a wall socket or to a vehicle's battery). Its true sinusoidal waveform is identical to household current. That means no special filtering or processing is required to power even the most sensitive electronic instruments. Unlike a conventional inverter, it does not introduce the unwanted switching transients associated with inductive loads which can introduce distortion into audio devices. And unlike a portable generator, it requires no fuel, makes no noise, and emits no internal combustion exhaust fumes. The FAR OUTLET Model 250S meets F.A.A. safety requirements for carrying aboard commercial airline flights."

For info on this product, check out this link: http://www.galaxyaudio.com/galaxy/galaxyPRODUCTS.html - PERSONAL POWER

The Galaxy Audio claims the Faroutlet will be available the second quarter of 2001.

This product would seem like a possible solution to the AC power problem. It remains to be seen how little noise is actually generated from this unit, and what the potential is for easy field use. The product weighs 31.5 pounds (which, like everything, has a way of getting heavier with every step). This might be an option for festival shows, but again, since 24-bit recording is capable of accurately capturing very low-level sonic information, it needs to be fully tested to see that it doesn't introduce noise or interference into our recordings. Again, the best bet is always to power your equipment with pure DC voltages.

1.3.6 Caring for your Equipment

1.3.6.1 How should I store my laptop while traveling?

A: Very carefully. Seriously, the screen is without a doubt the most likely component to break. Improperly packed laptops can arrive at its destination requiring costly repairs. It should be noted that the smaller the screen, the less likely it is to break. It is probably somewhat cheaper to replace a smaller screen than a larger one. A broken Sony Vaio Picturebook C1VN screen costs $1000 to replace. Still, packed carefully and properly protected, your laptop screen should arrive intact. Store your laptop preferably in a protective laptop bag or case. If you want to pack your laptop in with other taping gear, do so in such a way as to avoid any incidental pressure on the screen side of the laptop. Put the screen against a large, padded, flat, stiff surface, like a thin hard-cover children's book covered with a small towel. The surface should be larger than the screen to avoid any uneven pressure against the screen, especially in the center of the screen. If you are flying to your destination, always carry your laptop with you. Disconnect all cables from the laptop before packing.

1.3.6.2 How should I store my laptop while in use?

A: If at all possible, keep your laptop out of, or on top of, your recording bag. Position your laptop in a way to avoid someone bumping into your input connectors like the connector on Digigram VXPocket breakout cable. Make sure the exhaust vents are not blocked. Cover your laptop with a plastic bag to avoid a soaking from an unexpected change is weather, or a less predictable beverage passing overhead.

1.3.6.3 How often should I defragment my dedicated audio partition?

A: If you get into the habit of recording your files and offloading them as soon as possible (deleting them from your laptop soon afterward), you will inherently avoid drive fragmentation. Defragmenting a large partition can take quite a long time. Often it is easier and faster to remove everything on your audio partition and put back only the files you need than it is to run a defragmentation operation.

The only time you might need to defragment your audio partition is if you do a lot of frequent editing (especially cut/paste/delete operations) on files that have resided on your laptop for extended periods. Make it a habit to keep your audio partition as clean as possible before recording new material, and you should never need to defragment your drive. If you really feel the need to defragment your drive, do it overnight.

1.3.6.4 My laptop gets pretty hot while in use. Is this bad?

A: All computers generate some heat, some more than others. This heat is not inherently bad, but too much of anything is not good. Ordinarily, this heat is not a problem. The trick is to manage heat by keeping good airflow around the unit. Laptops have vents that exhaust hot air from the CPU, hard drive, and other heat generating components. Keep these vents unblocked at all times. Also try to keep the laptop away from other heat generating components in your recording rig, and definitely keep it out of direct sunlight.

1.3.6.5 What can I do to keep my laptop cooler?

A: Make sure there is good ventilation around the unit, and that the exhaust vents are unblocked. A white sheet or towel over the laptop in direct sunlight will reflect heat away from it. A cold pack wrapped in paper towels and placed under the laptop will help cool things off, but watch out for excessive condensation that may develop from the use of cold packs in humid climates. Honestly, with these basic techniques, your laptop should get no hotter than your DAT deck ever did.

1.3.6.6 Is it safe to operate a laptop or hard drive recorder in a humid climate?

A: Yes, as long as the humidity does not cause condensation to build up on the unit. Fortunately, unlike DAT decks, there are no "dew" lights that might come on at inopportune times. If humidity is still a concern, a desiccant like "Silica-Gel," placed in a plastic bag along with your laptop should help, but be mindful of the plastic bag's affect on heat dissipation.

1.3.6.7 What do I do if it's raining at an outdoor show?

A: Do the same thing you do for your DAT deck. Either don't record or keep it in a plastic bag and touch it as little as possible. Start your laptop early before the show at a time when it is not raining so hard.

1.3.7 Miscellaneous

1.3.7.1 What are the most common components to break on a laptop rig?

A: The screen is especially prone to damage while traveling. Be careful of the connectors sticking out of the machine as well (Firewire, PC Card dongles, etc.).

1.3.7.2 How does the laptop recording experience differ from the more conventional 16-bit recording experience?

A: The recordist must be a bit more aware and focused until they are experienced with the laptop, its software, and troubleshooting basic problems. However, once everything is set up and running, aside from managing battery life and the 2 GB wave file size limitation of most software, recording to the laptop can be a more relaxing experience. Levels don't have to be watched as carefully with 24-bit recording, you don't have to worry about repacking tape, running 1 minute leaders, and placing StartID's. If you're patching 24-bit out of someone's rig, you can be assured a place at the front of the chain(assuming you have a patch bay, or you come directly out of an AD2K+), placing you at less risk of problems ahead of you. Excessive sunlight can be a small problem, and heat is as much a concern as it is for a DAT deck. Be prepared to answer a lot of questions from curious onlookers(refer them to this FAQ if you like).

1.3.7.3 What are the biggest caveats of field recording with a laptop?

A: Believe it or not, accidentally hitting the stop button while recording has topped the list. This can easily be prevented by moving the mouse away from the record/stop button on the screen. Direct sunlight makes it more difficult to see the screen. Computers definitely don't enjoy beer as much as most people do. The connector on the Digigram breakout cable isn't all that robust, and needs considerable care when handling the machine. And again, be prepared to answer a lot of questions from curious onlookers.

1.3.7.4 I'm still not sure I trust laptops as a sole method of recording. Can I run a DAT simultaneously?

A: Only the Benckmark AD2K+ has simultaneous 24 and 16-bit output on two separate digital channels. This will allow you to record 24 bits to the hard drive and 16 bits with proper word length reduction to DAT at the same time. See 1.1.2.3 and 1.1.2.5 for a better description of this setup.

1.3.7.5 What are some other benefits of laptop recording?

A: You can see what you recorded almost immediately after you recorded it. You never waste tape. You can track the file out easier for CD's. There are no tape heads to get worn out. You can upgrade your rig without "throwing the baby out with the bathwater." If a better soundcard comes out, you can replace that component if you like. If your company gives you a newer laptop, you can just replace your laptop. If you wanted to upgrade a single component in a self-contained professional-grade hard drive recorder, you'd be out a lot of cash. See also 6.4.4.

2 Playback

2.1 Digital to Analog Converters

2.1.1 What kind of hardware do I need to maximize the benefits of my 24 bit recordings upon playback?

A: In theory, 24-bit resolution should offer 144dB of dynamic range. However, in reality, 144dB is well beyond the limits of today's analog circuitry, which optimally is capable of achieving about 120dB of dynamic range before noise and distortion take over the signal. Therefore, 120dB of dynamic range is considered the practical limit today. Unfortunately, most 24-bit soundcards do not have an analog output stage approaching 120dB of dynamic range. This is probably because of several factors. Soundcards in computers are very susceptible to the induction of noise by the emission of electromagnetic fields created by such components as switching power supplies, video cards, fans, and the motherboard itself. Even if a soundcard appears to be well shielded, there is only so much interference this shielding is capable of preventing. Low-level noise and interference generated by a computer environment can easily negate or mask the benefits achieved by a high-quality 24-bit A/D conversion process. Specifications for the analog stages of soundcards reaching the 112dB mark are considered to be very good

Furthermore, in order to keep a 24-bit soundcard (with digital and analog inputs and outputs) in the moderately priced range of about $500, lesser quality components are used on the cards than can be found in dedicated D/A converters costing the same. Therefore, the best way to hear the benefits of a 24-bit recording is to use a dedicated outboard D/A converter. Additionally, since most work in the 24-bit realm is taking place in the controlled environment of a studio today, soundcard manufacturers probably assume that an outboard D/A would be used anyway. Still, recordists on a low budget may find their 24-bit soundcard to be sufficient in sound quality to justify the efforts of 24-bit recording. After all, even with 112dB signal-to-noise ratio (about equal to a 19 bit recording), 24-bit recordists can enjoy an additional and very significant 20dB of dynamic range over and above that of their DAT recording counterparts.

2.1.2 What brands and models of outboard digital to analog converters are available?

A: There are a number of 20 and 24-bit D/A converters available, priced anywhere from about $500 for a decent, basic converter, all the way up to $10,000 and beyond for multi-channel DSP converters with upsampling algorithms intended to make lesser quality/resolution recordings sound arguably better. The Lucid DA9624 (http://www.lucidaudio.com/product_ad_9624_da_9624_info.htm) has been said to be a great performer (>114dB Dynamic Range, unknown S/N ratio & THD+N) and value in its price range of about $500. The Lucid was recently nominated for a 2001 TEC Award. It supports the three most common two channel digital formats (Coax, Toslink, and AES/EBU), and has separate gain controls for headphones (through a ¼" stereo jack on the front) and main outputs. It also has very nice, fast 20 segment LED level meters. Other reasonably priced common D/A solutions include the Midiman SuperDAC 2496 (http://www.midiman.net/m-audio/html/products/supdac.htm) with specs quoting 115.5dB of Dynamic Range, 117dB S/N ratio, and THD+N @ 0dBFS = -100dB A-weighted, lists for $299. While the SuperDAC 2496 does support all three digital input formats as well, it lacks the additional features of the Lucid, including headphone out, gain controls (of any kind), and peak level meters. Both the Lucid DA9624 and the Midiman SuperDAC 2496 support sampling rates from 32kHz to 96kHz. Pricier D/A converters include the Apogee PSX-100 (a 96/24 A/D and D/A in one unit). Where no budget limitations exist, Theta Digital makes some very fine D/A converters and DSP processors.

2.1.3 Are there any portable, 24-bit outboard D/A converters available?

As of May 2001, there does not appear to be any easy solution targeted to people looking for portable, battery-powered outboard 24-bit D/A solutions. However, the possibility of a companion D/A unit for a very popular A/D unit is something we could all hope for. Currently, the SuperDAC2496 requires a very simple 9V power supply, and looks like a simple 12VDC to 9VDC converter would work nicely. However, an analog output rig would have to be made since it doesn't have a headphone jack. The Lucid DA9624 comes with a wall-wart supplying the unit with 10.5VAC and two additional lower power supplies of 17.25VAC. The other option is the use of a 12V power inverter to supply the wall-wart with the 120VAC it requires. It is certainly portable, but without some kind of a power rig, is incapable of remote field use. Generally though, the soundcard in a laptop would be more than sufficient for basic monitoring purposes required in the field.

2.2 Software

2.2.1 Do I really have to install a $400 software package on my computer just to play back my 24-bit recordings?

A: No. For the PC platform, WinAmp (free for download from www.nullsoft.com) versions 2.75 support playback of 24-bit wave files. Additionally, 24-bit wave files compressed by Monkey's Audio Compressor can currently be played back in their compressed state in a crudely truncated and rounded 16-bit format by the Monkeys Audio WinAmp plug-in, which by all means can still sound very good. Of course, neither option reveals that glory of the 24-bit recording like going the distance to do it right, as discussed above.

2.3 Soundcards

2.3.1 What are some available soundcards for my home PC that supports 24-bit playback with S/PDIF, AES, or optical digital I/O?

A: Digital Audio Labs makes the CardDeluxe. http://www.digitalaudio.com/carddeluxe/default.htm

SEK'D makes the Prodif Plus, Prodif 96 Pro, and the Arc88. See www.sekd.com for more details.

Mark of the Unicorn(MOTU) makes the PCI 324 card and the 2408mkII interface. They also have a Firewire product called the MOTU 828. See www.motu.com for more details.

Steinberg makes the Nuendo 96/52 Audio card. See http://www.nuendo.com/hardware/hardware03.html fore more details.

Soundscape Digital makes the Mixtreme card. It is a 16 channel TDIF card generally for use with DTRS machines, but with the additional S/PDIF option, this card have proven to be very stable. See http://www.soundscape-digital.com/Products/Mixtreme/Intro/ for more info.

RME Audio makes the ADI and Digi96 series soundcards and converters. See www.rme-audio.com for more details.

In addition to the beloved VXPocket V2 soundcard, Digigram also makes the VX222, PCX924, PCX22, PCX820np, PCX822np, and the VX820 soundcards. See www.digigram.com for more details.

Midiman makes a host of soundcards that support 24/48 and 24/96 recording. See http://www.midiman.net/m-audio/html/digtlcds.htm for a list of their products, including the Delta 1010.

Crest Audio makes the FB-88 multichannel Firewire soundcard, due for release in "Spring 2001". See http://www.crestaudio.com/html/fb88.html for their product web page.

MH Labs also is coming out (soon) with a new Firewire-based multichannel product that can be used at home or in the field, called the Multi I/O 2882. If this product meets all of its promises, it is sure to be a winner in the 24/96kHz soundcard arena. See www.mhlabs.com for their website.

3 Editing

3.1 Processing audio

3.1.1 What kind of impact does the application of digital processing have on sound quality for 24-bit files?

A: Much lower impact than a 16-bit process on a 16-bit wave file. More resolution is retained with each process. As a result of most processes, extra bits of precision are required to accurately represent the audio. When working with 16-bit wave files, these extra bits are discarded in order to keep the file at 16 bits. The cumulative error from repeatedly generating extra bits of precision from each process and then discarding them by rounding is what can cause the resolution of lower level audio information to degrade rapidly. Retaining more information with every process reduces quantization noise generated from these types of rounding errors. As a layman's example, suppose I could only represent a number as an integer. Let's take the number 3. If I run a process which multiplies 3 by 1.25, and store the result as a rounded integer, 3*1.25=3.75 which gets rounded up to 4. Now lets say we run a second process that multiplies this result by 1.376234. If we do this to 3.75, we get 5.1608775. But if we do it to the previously rounded number 4, we get 5.504936, which we should round to 6. As you can see, there is a big difference between actual resultant with precision (5.1608775) and the resultant obtained through repetitive rounding(6). Of course, this a little bit (no pun intended) of an oversimplification, but you get the idea. The more precision retained after each process, the more accurate the sound is afterward. Processing in the 24-bit domain allows far more information to be retained with each process. It does follow that good 16-bit mastering techniques being employed today by knowledgeable digital audio engineers have a 16-bit file being converted to 24 bits before processing. All processing occurs in the 24-bit realm, and as a last step, the 24-bit file is reduced to 16 bits again with dither (word length reduction). The results of this methodology are not quite as good as starting out with a 24-bit file, but it is still a major improvement in accuracy over processing 16-bit files with 16-bit processes truncating and/or rounding repeatedly. This technique has often been called "upsampling before processing," and efforts to retain better precision may also have the file converted to a much higher sampling rate before processing as well. Converting a 16bit/48kHz wave to 32 bit/192kHz, processing, and converting back with dither is an extreme example of "upsampling before processing."

3.1.2 When should I process audio, before or after converting to 16-bit?

A: See section 3.1.1 for an explanation of the benefits of audio processing in the 24-bit domain. All audio should be processed in the 24-bit domain, or higher, and a word length reduction to 16 bits (if that's really what you want) should be performed with dither as the very last process ever to be performed on the audio file.

3.1.3 What types of processing should I perform?

A: General theory of mastering digital audio states that less is more. However, audio processing by experienced audio professionals can certainly yield a better sounding recording. For the purposes of the purist concert recordist/archivist, a basic normalization which increases levels (peak normalization, without clipping) can have the effect of "turning up the volume," which can easily give the somewhat false impression of a better sound. Compression and limiting of any kind should be avoided or used very sparingly so as to retain as much of the original dynamic range as possible. Any kind of processing in the 24-bit domain will be more accurate in the end, but this doesn't necessarily make the result of every process sound natural. Still, minimally, if you're looking to burn 16-bit audio CD's from your 24-bit wave file, you should consider a basic normalization process to maximize levels (and therefore resolution) before converting to16 bits.

3.1.4 Does it help audio quality to convert the file to 32-bit before processing?

A: Just as processing 16-bit audio in a 24-bit environment reduces rounding and quantization errors thereby retaining sonic accuracy, editing 24 bits files in a 32-bit floating point system is equally beneficial. Remember, the utilization of a higher word-length editing system does not improve the sound quality of the source, it merely assures that the best possible quality is retained throughout processing.

3.2 Word Length Reduction(WLR)

3.2.1 Is it hard to convert to 16 bits for CD's?

A: No. There are two simple methods. Most commonly, this can be done in any 24-bit editing software like Sound Forge 5.0. Performing word length reduction in software works well for folks utilizing a computer-based CD burning process. If you have a standalone CD burner, some software will let you play back the 24-bit file while performing a word length reduction to 16 bits with dither in real-time, which can then be fed directly into your stand-alone CD burner. Digitally passing the 24-bit output from a soundcard through a Benchmark SONIC AD2K+ will also allow word length reduction to occur externally, with 7 different dither/noise shaping algorithms to choose from. Output from the AD2K+ can then be fed into a stand-alone CD burner. One major advantage of the latter approach is that the mastering engineer can listen the final mix through the AD2K+ and determine what word length reduction formula best represents the original 24-bit mix.

3.2.2 What would happen if I simply feed a 24-bit signal into a standalone CD burner?

A: The CD burner will only accept the 16 most significant bits, and will discard the other 8 bits. This is called truncation, and can yield audible low-level distortion and artifacts due to the lack of dither.

3.2.3 Are there any hardware devices that will do WLR?

A: Yes, the Benchmark Sonic AD2K+ analog to digital converter has a feature allowing digital pass-through and conversion of a 24-bit digital input signal to a 16 bit, dithered/noise-shaped output signal. For 16-bit compatibility, it has 7 different dither/noise shaping algorithms to choose from, including TPDF(white noise dither without noise shaping), Noise Shaping, and “Near Nyquist,” similar to that used by Apogee’s converters with UV22.

3.2.4 If I intend to convert from 48kHz sampling rate to 44.1kHz for CD’s, should I do it before or after performing a WLR?

A: All processing, including sample rate conversion processing, should be conducted in the 24-bit realm, and word length reduction with dither and noise shaping should always be performed as a last step before burning.

3.3 Dither and Noise Shaping

3.3.1 When should I use dither?

A: Any time a process is run which gives a resultant requiring greater precision to store than the ultimate target format will allow, dither should be applied to smooth out the artifacts from quantization noise generated by the loss of precision. Word length reduction from 24 bits to 20 or 16 bits requires a loss of precision, and therefore, dither is needed to retain as much of the original sound as possible.

3.3.2 What type of dither should I apply to the audio?

A: The right answer is always to trust your ears. However, when using the AD2K+ to perform the word length reduction, 16NN2 and 16NS3 dither settings have been used with good results on typical rock and roll recordings. Additionally, some dither settings have less impact than others. Lesser impact dither settings should be used if there is any possibility that someone might run another process on the audio after dithering. Because dithering is a finalizing or finishing process, no processing should ever occur after a dither. However, there may be times you’d expect someone might run another process on it unknowningly, and it is these times that lesser impact dither settings should be used. Dithering algorithms work by performing mathematical operations on sonically important least significant bit(s). The result is that the low order bit(s) become somewhat randomized according to a probability function or curve. While this randomization of low order bit(s) is considered noise by itself, it has the effect of audibly masking quantization noise inherent in the PCM format, making the music sound more transparent, and allowing you to hear some sonically important information that would otherwise get lost in the noise floor. Applying dither a second time will hurt the quality of the recording, because as it adds and compounds noise produced by previously randomized low-order bits, thus raising the noise floor and decreasing dynamic range and low level resolution. Most people are not aware of this, as they frequently process (most typically normalization and sample-rate conversion) audio recorded utilizing an Apogee A/D converter with UV22 enabled.

3.3.3 What type of noise shaping should I be applying to the audio?

A: Generally, the choice of both dither and noise shaping is a matter of the taste of the listener, and is also dependent on the content of a recording. Some dither/noise-shaping algorithms sound better on rock-and-roll, and some may be better suited to classical or other sound sources. The best answer is to trust your ears and try different algorithms out to see what sounds best. Again, if you plan on processing the audio further, do not use dither and noise shaping, or use minimal impact dither and noise shaping algorithms. Also, in many implementations, a choice of noise shaping curves separate from a choice of dither is not an option. Apogee's UV22 is an example of this, as they combine their proprietary dither and noise shaping algorithms into a single option. The Waves L1-Ultramaximizer (software DirectX plugin) does allow dither and noise shaping to be chosen separately.

4 Storage Management

4.1 Compression

4.1.1 How can I maximize storage capacity?

A: You can use lossless compression programs to compress your audio data to make room for other recordings.

4.1.2 I've always heard that compression is bad, which is the reason MiniDisc doesn't sound as good as DAT. Is this not true?

A: There are several different kinds of compression. There is audio dynamic range compression, which makes the difference between low level sounds and high levels less significant. This type of compression is used to raise the level of quieter material without simultaneously raising the level of louder material. Limiters are a form of dynamic range compression.

Then there are the lossy audio compression algorithms like ATRAC and MP3, which save space by discarding sonic information least likely to be heard by the human ear. These compression algorithms take advantage of psycho-acoustic properties of audio (such as concept of frequency masking) as they relate to the perception capabilities of the human ear. These are the types of compression employed by MiniDisc, Digital compact cassette (remember those?), and Dolby Surround 5.1.

Finally, there is the type of lossless compression used to reduce the size of data stored in a computer. PKZIP, WinZip, WinRAR, WinAce, Stacker, LHA, Shorten, etc., are all programs that compress data losslessly by identifying redundant strings of information and using symbols to represent those strings. These programs are based on data compression algorithms like Huffman coding, Rice coding, and Lempel-Zev coding. Data compressed by these programs can be later decompressed back into its exact original form. This is the type of compression that can safely be used to save space on your hard drives, and has been in use for years in the computer industry.

4.1.3 What are my options for lossless compression?

A: To compress audio data, you can use general, non-data-specific compression programs like WinZip, or you can use compression programs like Shorten and Monkeys Audio Compressor (MAC) which are specially tuned to provide maximum compression specifically for audio data. Programs not made specifically for compressing audio data will not achieve nearly as high a compression ratio as programs like Shorten and MAC.

4.1.3.1 Almost everyone on Etree uses the Shorten wave compression format. Does Shorten work on 24-bit wave files?

A: No, unfortunately not at the moment. Only Monkeys Audio Compressor and the RKAU audio compressor allow lossless 24-bit wave-specific file compression. The following URL links to a comparison chart to see where these and other lossless audio compressors stack up.

4.1.3.2 What is Monkeys Audio Compressor?

A: Monkeys Audio Compressor is a freely available, closed-source, Windows-only, wave-specific lossless wave file compression program written and maintained as a hobby by Matthew Ashland. It offers significantly faster and tighter compression than Shorten, and decompresses faster as well. Furthermore, unlike Shorten, it supports the compression of 24-bit wave files. A WinAmp 2.0 plugin is available to play Monkeys Audio files (.APE files) in their native format. A plugin is also available for Media Jukebox 3.0. See www.monkeysaudio.com for more information on this fine program.

4.1.3.3 If Monkeys Audio Compressor (MAC) is faster and offers better compression than Shorten, why isn't everyone using it instead of Shorten?

A: The most major reason is that versions of the Shorten program and its algorithms are also available for non-Windows platforms like Linux and Mac OS. The author of MAC has recently been making some strides in improving the design of his program to make it easier to port to other platforms like Linux. However, since the program has always been targeted to Windows users, and it is just a hobby for Matt(a.k.a. Monkey), making cross-platform versions of MAC is not a high priority for the author. Matt has also decided to keep his program "closed-source for the time being, as is his prerogative. In any case, Matt's programming has given the pioneering 24-bit field recordist a great option.

4.1.3.4 Can't I just use WinZip or WinRAR or WinAce to compress the files instead?

A: You certainly can. However, because these programs are not tuned specifically to the rather random data of audio, it will take far longer to compress your files, and you will not gain much. You'd be lucky to squeeze a 1.5GB 24-bit wave file to 95% of its original size if you used WinZip to compress the file. It would probably take over an hour on an average PC to do it as well, and you can expect decompression to take a similarly unreasonable amount of time. Compare this to Monkeys Audio Compressor, which will compress an average 1.5GB file to 60-70% of its original size in less than 10 minutes. As you can see, it's really not much of a question. Aside from using MAC, WinAce can achieve the best compression (consistently 77%) of a 24-bit wave file of all the standard data compression programs. However, WinAce will take hours to perform this operation (and a long time to decompress as well), while MAC will deliver significantly better compression and is much faster to boot.

4.1.4 Should I be concerned about the availability and reliability of Monkeys Audio Compressor (MAC) in the future?

A: Matt Ashland, the author of MAC, is committed to maintaining backward compatibility with previous versions of his software. That said, he makes no guarantees as such. The best guaranty of being able to decompress an .APE file into its original .WAV format is to archive a copy of the version of MAC used to compress the file along with the file itself. Fortunately, MAC is a very small, lightweight, easy to install program.

4.1.5 I like to run Linux, but Monkeys Audio Compressor only works on Windows. What can I do?

A: For Linux users, MAC has been successfully run under Windows emulators such as WINE.

4.1.6 How can I verify that I definitely haven't lost anything after compression?

A: Monkeys Audio compressor contains a verify option that can ensure that your file can be decompressed with bit-for-bit accuracy.

4.1.7 Can I play the compressed files back without decompressing first?

A: Yes, there are plugins available for Nullsoft's WinAmp and Media Jukebox to be able to play .APE files without decompressing. However, because of previous limitations in WinAmp's ability play only 16-bit files, the MAC WinAmp Plugin detects 24-bit files and forces you to use the "scale output to 16 bits" before the file will play in WinAmp. The newest version of WinAmp no longer has this 16-bit maximum limitation, but the MAC plugin has yet to make "scale output to 16 bit" a mere option. Matt Ashland, author of MAC, is now aware of this, and he is looking into modifying his plugin to better support us 24-bit folk.

4.2 File Management

4.2.1 Is there a limit to the file size of a wave file?

A: Yes, the limit is 2GB. This equates to about 2 hours of 24-bit 48kHz recording.

4.2.2 Is there any way around the file size limit?

A: This is a limitation of the .WAV file format. Some programs, however, attempt to provide a solution to this with a proprietary, program-specific file format. One such program is Sonic Foundry Vegas Audio/Video 2.0d, which allows wave files larger than 2GB to be saved after recording as .W64 files. These files can only be read by Vegas Audio/Video 2.0d, but Vegas can be used to split the giant .W64 file up into smaller .WAV files less than 2GB in size. The best thing you can do is to prevent any single file from ever going beyond the 2GB mark. If you know you will need to record continuously for longer than 2 hours, stop recording and start recording a new file as quickly as you can at a convenient moment between songs. This can be easily done while recording with SoundForge 5.0 with two mouse clicks inside of a second. Attempting to record beyond the 2GB with most software may result in a complete loss of your recording. This is something that the recordist should test out.

4.2.3 Should I look into other file formats like AIFF?

A: You can, but they too will have the same 2GB limitation.

4.2.4 If at some point, I plan to make audio CD's from my recording, should I break the file up into smaller files or tracks?

A: You should consider the following information before deciding which storage route to take. First, if you are not recording long sections of audio (over 45 minutes without a break), you may as well store your files in their uncompressed form (i.e. .WAV files). If, on the other hand, you are recording material that is not conducive to easily splitting up over multiple volumes, you will find it quite useful to apply a lossless compressor and/or an archiving system. If you prefer to avoid the use of an audio compression or archiving tool, you will need to cut the uncompressed .wav files into sections that can be reassembled seamlessly. While splices between tracks are typically easy to restore, recording bands that segue many songs without stopping can prove difficult. Using a CD that only allows for 650-700MB will only provide room for 45-50 minutes of audio at 24 bit/48KHz. Choose carefully!

Keep in mind that when viable 24-bit playback mediums become widely available, you may feel differently about how you previously tracked your recordings, based on things like the maximum amount of playback time available per media unit. The common 74-80 minute CD barrier is sure to be history as DVD-Audio becomes more ubiquitous.

4.3 Hard Drives

4.3.1 What drives provide the best performance when working with 24-bit files?

A: Ultra2 and Ultra160 SCSI-3 Wide 10,000 RPM hard drives are the fastest drives on the market. Using more than one of these drives in conjunction with a RAID controller will give you performance that meets the demands of things like video editing. This solution is standard for all server-class machines found in the business world, like large database servers, file servers, and web servers. For performance, it can be generally stated that SCSI drives are faster than IDE/ATA drives, although IDE based RAID solutions are becoming more common and rival the performance of the more expensive SCSI RAID solutions.

4.3.2 What drives are the most cost effective?

A: IDE and ATA-100 drives in sizes just below the largest cutting-edge drive size typically offer the best cost per gigabyte ratios. If the largest IDE drives available today are 80GB, the best deals can probably be had on 60GB drives. Also, slower drives are always cheaper than faster drives of the same form factor. If you plan to use a drive for archival of 24-bit files, the speed of the drive is not as important. IDE RAID array controllers are readily available today to help you build a decent performance audio workstation using the cheaper IDE/ATA drives.

4.3.3 I can get a good deal on an 80GB hard drive. How many should I buy to get started?

A: Because hard drives are constantly and very rapidly falling in price and growing in capacity, you should never buy more than what you need at any given moment, unless you are willing to pay the extra price to have free space instantly on hand in an emergency. A month after you find a good deal on a drive, a significantly better deal will become available.

4.3.4 How often should I defragment my drives?

A: See section 1.3.6.3.

4.3.5 How should I partition my drives for best performance when working with 24-bit files?

A: Partition your hard drive so that audio data is completely separate from your operating system and software. To decrease the need for defragmenting, keep your audio data partition sizes at 8GB or less where possible. While this limits your continuous record time to eight hours for 24-bit/48kHz, we are confident that you should be able to find at least one stopping spot during such a session in order to select the next drive partition for further storage. One big 16GB partition is much more likely to become fragmented than two smaller 8GB partitions. Also, the smaller the partition, the less time it takes to defragment it. Additionally, the audio partition should be an extended, logical partition. Do not create your audio data partition as a primary partition, or use the dynamic partition type offered by Windows 2000. Use extended/logical partitions instead. If you have any problems with your machine and you need to rebuild the OS, or you need to get the data off your machine from another OS installation, it may be difficult or impossible to recover your audio data if you use dynamic or primary partitions. Dynamic and secondary primary partitions are unreadable from a dual-booted installation of Windows 98/ME.

4.3.6 How should I format my partitions?

A: PC users should format their partitions using FAT32, not NTFS. Again, if something goes wrong with your PC and you need to rebuild the OS, you risk losing access to your audio data partition if you use the NTFS file system. NTFS formatted partitions cannot easily be read from any other platform other than Windows NT/2000.

4.3.7 What are my options for removable hard drive solutions?

A: There are several options that work well specifically for laptops, and others that are designed for a desktop machine. Laptops always have at least one PCMCIA, PC Card, or cardbus slot. This allows any number of drives and interfaces to be connected to your laptop. A PCMCIA IDE/ATA card will allow you to buy an external IDE/ATA chassis and a large IDE/ATA drive. You can do the same thing with a PCMCIA SCSI card and an external pullout SCSI hard drive enclosure, but with this solution, you must be sure to buy the less powerful 50 pin Ultra SCSI-2 drives, because these cards (like the Adaptec AHA1480 CardBus card) do not support the 68pin or 80 pin SCSI-3. See www.adaptec.com for more on that product. If your lucky enough to have one of the new laptops with IEEE1394 FireWire or iLInk400 ports built in, or if you purchase a PCMCIA FireWire adapter, you can get a portable IDE to FireWire chassis ($125.00) and a 60GB drive for about $180.00. Both of these solutions require external power for the hard drive chassis. USB hard drive solutions should be avoided because of their significantly slower speeds (1.5MB/sec maximum transfer rate), in addition to the fact that USB is quickly becoming eclipsed by the superior FireWire ports (50MB/sec maximum transfer rate) available on newer laptops. If you add a FireWire card to your home PC for about

$50-60, you can quickly offload a show from the laptop to the external FireWire drive, detach the drive from the laptop, and connect it to your home PC in less than 10 minutes.

4.3.8 Should I consider removable hard drive solutions for temporary storage?

A: Absolutely. A 60GB IDE/ATA100 bare hard drive that can be placed in a $125 IDE to FireWire chassis, costs about $180.00. This means it is possible to store 24-bit/48kHz recordings for as little as $3 per recorded hour. For the concert tape collector, this is only marginally more than what you would spend for 2 60meter DAT tapes. If you and your friends are all 24-bit tapers utilizing the IDE to FireWire solution, you can trade shows at a rate of about 8 shows an hour in the parking lot of your favorite venue with a 12V power inverter in your car's cigarette lighter providing 120VAC for your hard drives. For the recording professional using audio grade CD media, or DVD-R/DVD-RAM, hard drives are actually more economical. Use of lossless compression can bring the hard drive cost per hour ratio down to about $2.

4.3.9 Should I consider removable hard drive solutions for more permanent storage of my files?

A: A hard drive bought today is OK for storage for about 5 years. Technically, the hard drive may still work longer than 5 years. However, if you think about the cutting edge 1GB hard drive you may have bought in 1995, and what it looks like today compared to 60 and 80GB drives, its not too hard to see that you risk for buying a storage medium that is soon to be very obsolete. IDE/ATA has been a hard drive standard for a while now, and things change rapidly in the computer industry. Its very questionable that the IDE/ATA interface of today will be around 5 years from now, and you may have trouble recovering 25 or more shows if you can't read the contents of a hard drive. Also, hard drives are electro-mechanical devices that can be damaged by shock and static electricity, and can have problems spinning up after sitting for a long time. Therefore, in addition to keeping your files on hard drives, you should also consider burning audio data CD's or some other form of storage medium as a backup. DVD-RAM discs can be had for about $10 a piece in spindles of 50, and as their cost continues to decline, they will probably be better suited for long term archival.

5 Archival and Trading Methodologies

5.1 CD-R

5.1.1 Is it possible to burn CD's from 24-bit audio files that I can play on a regular CD player?

A: Yes, but only if you convert the 24-bit files to 16-bit files and master an audio CD. There is no way to make a CD-R with 24-bit native audio that will play in a regular CD player. The closest thing that comes to this is the Alesis Masterlink CD 24-bit CD recorder.

5.1.2 How many CD's will it take to archive an average 24 bit/48kHz show if I compress it with Monkey's Audio Compressor first?

A: Depending on the length of the show, anywhere between 3 and 5 discs. A 2.5 hour show can often fit on 3 74-minute CD's, and a 3.5 hour show sometimes requires 5 CD's. 80 minute discs may slightly reduce the overall number of discs needed for each file to be stored. However, the author doesn't believe that the arguably riskier 80 minute CD's are worth the savings of a very occasional $0.30 CD. Overburning a 74-minute CD an extra minute or two usually gets you by if it's close. Increasing the compression level on MAC can also give you that small edge you might need, but it does come at the cost of speed of decompression, and you will probably have to perform the additional step of recompressing the file again. Fortunately, Monkey's Audio Compressor allows changing the compression ratio of an .APE file in one step. While it may seem petty to be concerned with the cost of a single additional disc, remember, it is an additional step to burn the disc and to load it for playback. In the long run, these are all relatively minor expenses in the scope of going the distance for 24-bit field recording.

5.1.3 Should I consider CD-R as a viable archival and trading format for 24-bit audio today?

A: Yes. Because they are data format CD's you can clone them very quickly without risk of errors with newer computer-based 8x and 16x CD burners. A 24x CD burner from Plextor is due out very shortly, which will allow you to copy a CD in 2 minutes. Yes, it takes a bit of work to set up and burn the initial CD's, but nothing beats burning a high-quality show for someone inside of 10 minutes. For the time being, because of a lack of easily accessible native 24-bit playback formats, 24-bit archival and trading is a bit of a labor of love. So CD-R burning, along with hard drive cloning and internet based trading are probably the most accessible ways to distribute your recordings today. Look for DVD-RAM and DVD-R to replace CD's in the near future.

5.1.4 If my files are much larger than a CD will allow, how can I break the files up on CD's?

A: You have several choices for archiving audio on CD's. If you don't want to compress your recordings with Monkey's Audio Compressor, you can break your wave files up into CD sized chunks by using your editing software. At 24bit, 48KHz, you can fit about 45 minutes of music on a CD. This will allow you to play your files directly off of the CD's using a computer in full 24-bit quality. If you choose to compress your audio with MAC, you can break your wave files up into 55-60 minute chunks before compressing, and the MAC compressed files should also fit on a CD. This will also allow you to play your files directly off the CD's using WinAmp with the MAC plugin, although with the current version of the MAC plugin will only allow you to play 24-bit files if you "scale output to 16 bits." WinAmp has only recently acquired the ability to play 24-bit files, and the MAC plugin has yet to be modified to allow all 24 bits to pass through to WinAmp. Still, you can preview your recordings at 16-bit directly from the CD this way. The last option for those who don't want to bother with chopping up the large audio files is to use a typical robust data compression program like WinAce Archiver 2.0 to create CD-sized multi-volume archives out of your large files. Using this method, you will have to restore all CD's worth of data to your hard drive before playing.

5.1.5 What is the best method combining both speed and compression I can use to archive my large audio data files easily to CD's, without messing with an editor?

A: Compress your files using MAC set to normal compression. Then use a combination of the SFX Factory and WinAce 2.0 to create a self-extracting, multi-volume archive of your .APE files. Set the compression in WinAce to "None" or "Store". Do not enable additional compression in WinAce, because your audio has already been compressed, and it will take WinAce an eternity to try to squeeze more out of your file (which, in the end, it can't). Set your volume size to 650MB CD size, and enable a Windows GUI based self-extracting header to be placed on the file. WinAce will let you put some information about the archive in a little window which will be displayed when you try to extract the files from the CD's, and it is a good idea to fill this information window in with items like file sizes, equipment used in the recording, etc. It should take about 10 minutes for WinAce to create the self-extracting multi-volume archive on your computer (make sure you have as much free space as the original compressed files require to store the output from WinAce). Your laptop will then have multiple files of 650MB or less on them that you can burn to CD's. One of these files will have an .EXE extension, and the rest will have extensions like .C00, .C01, .C02, etc. Burn the .EXE file to disk #1, .C00 to disk #2, etc. To extract the files from the CD's to your hard drive, simply insert CD #1, click on the .EXE file, and insert the next CD's when prompted. It is a good idea to put a copy of the version of Monkey's Audio Compressor program used to compress the audio on at least one of the CD's, to ensure the ability to decompress your files in the future. A small plain-text "description.txt" file placed on the CD will also help identify your archive set. See www.winace.com for details on the WinAce Archiver and the SFX Factory, which are available for download for a nominal fee.

5.2 DVD-R

5.2.1 Where is DVD-R technology today?

A: DVD-R drives are capable of burning discs that are readable in your average home DVD player. DVD-R Recorders from Panasonic are available for about $3000. According to one large supplier of blank CD and DVD media, blank DVD-R's can be had in spindles of 50 for about $10 a piece. While this may be an affordable solution for some, it does not yet equate to a very cheap solution for the general population's need for 24-bit audio archival. The availability of DVD-Audio mastering software is currently fairly limited as well.

5.3 DVD-RAM

5.3.1 Where is DVD-RAM technology today?

A: DVD-RAM drives are capable of burning discs that are only readable by another computer-based DVD-RAM player, or by some of the very latest (and most expensive) stand-alone DVD players. As time goes on, we can expect greater compatibility of stand-alone players with all types of DVD media, as well as cross-compatibility for CD's and SACD's. DVD-RAM Recorders are available for about $500. Blank DVD-RAM disks can be had in spindles of 50 for about $10 a piece. While this may be an affordable solution for some, it does not yet equate to a very cheap solution for the general population's need for 24-bit audio archival.

5.4 Hard drives

5.4.1 Can I use removable hard drives like external FireWire drives for sharing my recordings?

A: Yes, and believe it or not, using hard drives as a trading medium provides the best balance between ease, speed, and cost. Trading 60MB hard drives, you can clone 20-25 shows in about 2 hours. It costs less than $10 to ship a 2lb. hard drive through the mail with insurance. You can trade a hard drive with or without the external chassis, although it helps for both parties to have exactly the same make and model of external chassis. See also section 4.3.9 for a discussion of long-term archival on hard drives. Some live touring acts are already utilizing such an approach for recording all of their performances. Of particular note, The Dave Matthews Band is presently making 48-track recordings of each show. Using duplicate sets of hard drives, they take all of these recordings on a weekly basis, send them to a safe location where the hard drives are backed up to CD's (about 60-80 per show) and then reformatted while the others are in use. The rotation of these drives makes it possible for them to continue documenting every performance in 24-bit splendor.

5.5 Network and Internet trading

5.5.1 Can I share 24 bit .WAV or .APE files over the internet like E-tree does?

A: Yes. It is only a matter of time before 24-bit FTP sites make material for download and trade over the internet. Known 24-bit recordists today do plan on setting up limited access FTP servers for downloading 24-bit audio recordings.

5.5.2 Can I exchange recordings with other laptops while on the road?

A: Yes, there are two good options for this. If you have a PCMCIA 100Mb Ethernet card or laptop comes equipped with a built in 100Mb Ethernet port, you can use a "crossover" cable to hook two laptops directly together via Ethernet. Crossover cables are usually available anywhere you can buy CAT5 UTP Ethernet cable, or you can make your own (test it fully first).

Another option for Ethernet (using a regular, non-crossover cable) is a small pocket-sized 100Mb Ethernet switch or hub, which can be had for less than $100. A hub or switch will let you share your music files with multiple laptops at once. However, keep in mind that the hard drive in your laptop may "thrash" if two people pull two different files from your laptop at the same time. This occurs because the heads in the hard drive would be constantly moving back and forth (called "seeking") between one file and another, rather than just reading data consecutively from one file. One of the specifications of a hard drive is called "average seek" time or latency, measured in milliseconds(ms). This number is the average time it takes to perform a head seek from one track on a disk to any another. A head seek is needed to reposition the heads over data that does not reside in the current physical track. No data can be read until the head seek is complete. As you'd expect, head seek time is directly proportional to the distance the heads have to travel. The fastest seek times are head movements from one track to the next, as is the case when the drive is reading a single, unfragmented file at a time. It is because of excessive latency caused by asking the hard drive to read from two completely sections on the disk at the same time that the hard drive heads "thrash" violently back and forth, which wastes time and causes the aggregate concurrent transfer of both files to take longer than the total time it would take to read the files consecutively.

In order to get your laptops talking to each other, you will have to configure your IP addresses on each machine so that they are on the same logical (as well as physical) network. For example, give laptop #1 an IP address of 10.0.0.1, and laptop #2 an IP address of 10.0.0.2. Set the subnet mask on both machines to 255.0.0.0. You can set your gateway address to the other machine's IP address as well. Then you can map network drives between the machines and copy files between the machines. A three hour file should take about 15-20 minutes to transfer using 100Mb Full Duplex Ethernet. Do not try to use 10Mb Ethernet, or USB, InfraRed, or parallel port file sharing solutions like Laplink, as they will be too slow for such large transfers.

The second option is only useful if both laptops have iLink400 or Firewire ports and both also have Windows ME installed. Sony Vaio Picturebook's come with Windows ME preinstalled, and if you installed Windows 2000 on the machine in a dual-boot setup (rather than completely wiping Windows ME off the machine), then you can use this second option. Windows ME supports networking two machines over a FireWire or iLink 400 cable. This will give you a high-bandwidth connection between the two machines, and again, you can copy files from one machine to another faster than Ethernet, if you have fast enough hard drives in your laptop.

5.6 High Resolution Compact Disc

5.6.1 Can I archive and trade 24-bit high resolution Compact Discs made on the Alesis Masterlink CD recorder?

A: Yes, but obviously, both parties must own an Alesis Masterlink CD recorder. Additionally, at about 45 minutes for a 74 minute CD, a basic 3 hour recording can easily require 5 CD's. You must chop up, track, and master your recordings for this limited format, which can be a big effort for such a temporary storage medium. The Masterlink CD recorder is much better suited to studio use for creating a quick, high-quality demo of a song or two than it is for live concert mastering. This makes this solution not the most viable of all trading mediums and methodologies.

5.7 High Resolution DTRS tape

5.7.1 Can I archive and trade 24-bit material using High Resolution DTRS tape?

A: Yes, both parties will need to own two Tascam DA78HR DTRS machines, or one DTRS machine and one 8 channel 24-bit TDIF soundcard like the SoundScape Mixtreme, the MOTU PCI 324 with a 2408 interface, or the Digital Audio Labs TDIF soundcard. The longest DTRS tapes allow 1h48m of total record time per track. Again, since these are 8 track machines, you can use tracks 1 & 2 for the first set of a performance and tracks 3 & 4 for the next. This potentially leaves 4 more tracks for another concert, meaning you can clone 2 shows at a time. The obvious benefits of trading and archiving 24-bit audio on DTRS tape is that you can listen back relatively easily without a computer. The caveats of this method are the facts that you are limited to real-time tape copying with two DTRS decks, or you have to bounce it to the computer first, and then out to another DTRS tape if you have the single DTRS deck and soundcard scenario. Also, DTRS tape stores audio in audio format similar to that of DAT tape, without all of the benefits (like longer lasting data integrity) in longevity and resiliency that data-based storage formats provide. This means you will probably have to reclone your DTRS tapes every 5 years, or they will deteriorate just like DAT tape. And of course, like DAT tape, the more you play them, the faster they deteriorate.

5.8 Digital Linear Tape

5.8.1 What is digital linear tape(DLT)?

A: Digital Linear Tape (DLT) is a tape backup format used in data centers around the world to backup and restore large amounts of data very quickly. It makes use of a technology called Linear Serpentine Recording, which stores data on 64 tracks. After each pass of the tape, the heads advance to the next track, and the tape reverses its travel, making a giant “S.” It is because of this technology that a particular data set can be very quickly located on tape. Tape in a DLT cartridge is wound upon a single reel. When inserted into the tape drive, the tape is immediately spooled onto a take-up reel inside the DLT drive, and the tape header is read. In its common 35/70GB format (known as DLT7000), DLT can achieve backup rates of up to 300-500MB/min. A DLT7000 tape drive requires DLT Type IV tape cartridge, costing from $25-65, depending on branding, packaging, where you get your tapes, and whether they are new or not. A DLT Type IV cartridge can store 35GB of uncompressed data, or 70GB of compressed data (the compression occurs in the hardware of the DLT drive). DLT tapes are said to have a shelf life of 30 years or more, and get better with use (believe it or not) because the tape becomes more frictionless by polishing that occurs as the tape flows across tape heads repeatedly. Used tapes are not recommended, even if they are so called “once used” tapes. DLT tapes provide the absolute best cost per megabyte over all trading methodologies. About 20 compressed average length shows on a $50 tape yields a cost of about $2.50 per show. Of course, this low cost can easily be offset by the cost of the DLT7000 drive itself. However, as data centers begin to phase in newer cutting edge tape technologies such as SuperDLT tape (offering 50/100GB or greater capacities), used 35/70 drives are sure to become a bargain. For more information on DLT technology, including proper care and use of DLT drives, visit http://www.dlttape.com

5.8.2 Where can I get a DLT drive for a reasonable price?

A: Many 35/70GB DLT tape drives have been posted on E-Bay, buyer beware. $4000 is about the most they would cost, and many times they can be had for $2500 or less. A DLT tape drive offers its owner the additional benefit of being able to back up their computer(s) easily and quickly.

5.8.3 Should I expect a DLT solution to ever become a native audio format?

A: No. DLT is marketed only as computer backup device, usually used to back up servers in data centers. While it would certainly be very nice for the audio industry to make use of DLT technology, the architecture doesn't lend itself easily to audio, so don't expect this to change.

5.8.4 Do I need to purchase backup software to use DLT tape?

A: Yes, a product like Veritas Backup Exec 8.5 for Windows NT/2000, or ArcserveIT for Windows NT/2000 (from Computer Associates) is needed. If you own a DLT drive, do not attempt to use the backup software that comes with your operating system in place of a full blown robust package like Backup Exec, or you will very much regret it later. The Veritas Backup Exec Product does require a server class OS like Windows 2000 Server. Veritas does make a home version of BackupExec which will work for Windows 2000 Professional. The ability to install the ArcserveIT product on Windows 2000 Professional is unknown. Registry tweaks on a Windows 2000 Professional installation may be able to be made to make Windows 2000 Professional report back to other programs that it is really Windows 2000 Server. Performing these tweaks may leave you in violation of Microsoft's copyright though, and should be done at your own risk. BackupExec has specifically been shown to be a great, stable, and reliable product for backup purposes, and is very widely used in data centers everywhere.

5.8.5 How can I use DLT tape for archiving and trading?

A: Each party in a trade will need at least one DLT tape drive. Two DLT drives will allow DLT tapes to be copied easily in one step. Otherwise, a DLT tape would have to be copied onto a hard disk, and then backed up onto another DLT tape. For archiving, DLT tape is not recommended unless the recordist is willing to keep two or even three copies of a given tape. DLT tape is generally reliable, but it is still tape, and if anything goes wrong with your only copy of 20 shows, you may be out of luck. DLT tape does make a great trading medium if you archive to external or pull-out hard drives, and trade using DLT. This way, you can connect your hard drive and make a DLT tape backup of the drive (and multiple shows) at a rate of about a show every 10 minutes. The recipient can choose to restore shows and archive them as he sees fit, either to hard drive, or to multiple CD's. As a greatly added benefit, the DLT tape can then be kept as a backup of the hard drive as well. Techniques such as making extra DLT backups of a hard drive in advance can also save a lot of time. DLT hardware compression should always be disabled when backing up audio data, because different DLT drives (or drive firmwares) may have trouble reading the compressed tapes from another drive. If compression of audio data is desired, it should be done by Monkey's Audio Compressor before backing up to DLT.

6 Miscellaneous

6.1 Equipment Manufacturers

6.2 Retailers

6.2.1 Where can I find retailers dealing in 24-bit audio products and solutions?

A: See Sonic Sense at www.sonicsense.com for 24-bit A/D and D/A converters (including the Benchmark Sonic AD2K+ A/D converter, the Lucid 9624 D/A converter, and Apogee products like the Rosetta or PSX-100), 24-bit soundcards (like the Digigram VXPocket V2, and the Sekd Prodif 2496), and high quality microphone preamps like the Grace Lunatec V2. Marc Nutter at Sonic Sense has been one of the pioneers of 24-bit audio field recording, making him very knowledgeable and supportive of 24-bit products and field recording solutions. While slightly better prices can occasionally be found at less knowledgeable vendors, many recordists find the value-add Sonic Sense provides its customers to be worth a bit more. NovaMusik at www.novamusik.com also sells 24-bit soundcards like the Soundscape Mixtreme TDIF PCI cards for some of the best prices on the net. They are also a good place to get the Midiman SuperDAC2496.

6.2.2 Where can I find turnkey 24-bit audio solutions?

A: See SonicSense at www.sonicsense.com .

6.3 Market Pricing

6.4 Other Resources

6.4.1 Where can I find other people who are doing 24-bit field recording?

A: Since the only affordable means of 24-bit field recording involves use of a laptop, you can find other pioneering recordists on the laptop-tapers@yahoo.com mailing list.

6.4.2 Are there any sites similar to Etree that are targeted towards 24-bit field recordists?

A: Not now, but expect them soon.

6.4.3 Where can I find more information on digital audio?

A: A great book explaining all facets of digital audio, including media formats, encoding schemes, and theory, is “Principles of Digital Audio,” by Ken Pohlmann. Look for it on Amazon.com or your favorite bookseller.

6.4.4 What are some other nifty things you can do with your laptop?

A:

1. You can get a wireless (cellular) internet card and service from companies like Ricochet, which provides a symmetrical 128Kbps of internet bandwidth for $99/month. This will allow you to browse the net from anywhere.

2. If your laptop has a built-in soundcard and a headphone jack, you can listen to your recordings on long trips directly from your laptop into your car stereo with one of those cassette adapters. You will definitely not get the type of quality you’d expect from a 24-bit recording, but you can still listen to your recordings on the road.

3. If you have a high-bandwidth internet connection at home(1.54Mbps or more downstream), like DSL or cable modem, you can set your machine up as an FTP server. Then, while you're on the road, if you're lucky enough to be staying at a hotel with Ethernet-based internet access, you can use FTP to transfer your recordings to your home machine. At 1.5Mbps, it will take about an hour to transfer 580MB, so plan on 5-7 hours to transfer a show to your home machine over the internet. This can be done overnight, or if you are lucky enough to have another laptop with you (perhaps an older, junkier one) with enough space on it for one recording, you can transfer your files from your recording laptop to the secondary laptop, and use the secondary laptop to send your files over the internet while you use your primary laptop for more recording. Perform a test with a small file and a stopwatch to see what kind of upstream speeds you achieve from the hotel room, before you attempt to send 3GB over the internet.

4. At festival shows, if you can get A/C power, or you have a car battery and a 12V power inverter, you can power your externally attached FireWire hard drive and either record directly to it, or you can copy files off your laptop onto the FireWire drive between performances.

5. You can play games on the plane or in the car while traveling. Favorites games played on laptops while traveling are MAME-based games. MAME is an arcade-game emulator, and you can download original arcade game ROM's from the net and play them without dropping a single quarter. Joust, Phoenix, and Millipede are amongst my personal favorites. For the violent folk with a 24-bit laptop recording companion, a nice head-to-head two player game of Doom or Quake can help you get out some aggression.

6. If your laptop has a USB port, you can get a USB hub, and a USB keyboard and mouse. Then you can attach your video output to a regular monitor, and you can pretend your laptop is a desktop machine while at home. This solution is much cheaper than a docking station or port replicator, although they would also work for the same thing. A docking station for a laptop usually offers the added ability to use full PCI cards(like multitrack soundcards and RAID or SCSI disk controllers) with your laptop.

7. You can watch movies in DVD, Video CD's or DIV-X formats while traveling.

8. You can purchase a wireless Ethernet (802.11) card for your laptop for about $150. These cards provide up to 11Mbps of bandwidth, depending on how far away you are from your nearest wireless access point or tower. Most major airports now offer wireless internet access for a nominal price($5-10). You simply insert your wireless card, boot up your machine, open up a web browser, and supply a credit card. Then, while you wait for your delayed flight, you can do all kinds of internet stuff like email and web browsing, including checking the status of your flight, or booking a new one. With a cheap microphone, you can even use services like Net2Phone to make inexpensive, long distance phone calls. Try that with your DAT deck!

9. You can master 16-bit audio CD's on the road with a portable CD burner.

10. With a laptop with a built-in camera, you can take digital pictures of your friends as they attempt to escape reality in the parking lot before the show.

7.0 The Future

7.1 Encoding Schemes

7.1.1 What is Direct Stream Digital(DSD)?

A: ***Coming soon.

7.1.2 What is the future of Pulse Code Modulation (PCM)?

A: Currently, the biggest buzzword in the audio industry is 24/96, but there is already a move afoot to step up sampling rates to 192kHz. We can easily see how big these files will become: 4 gigabytes per hour for stereo. Despite the massive file size, some producers, like Elliot Mazer (Neil Young), are already utilizing 24/192. "I am sure we will see great products such as Pro Tools, SampleCell and Vision working at 192kHz/24-bit. When this happens, these tools will let us create recordings that sound analog, and to me, that is a good thing." MIX Magazine, Chris Michie, Vol. 25, No. 5

How far PCM will be taken is yet to be seen. There are still professionals who argue that the difference between 16-bit and 24-bit is more valuable than the difference of 48kHz versus 96kHz and that the space required for 96kHz (or higher) is not justifiable for the benefits it provides. Nonetheless, those driven by ultimate sound quality, without regard to practicality in the present, may push technology further yet. One thing for sure is that 24-bit simply proves that there is still further we can all go with digital representation of the ever present analog sounds we hear.

7.1.3 What is Dolby Surround 5.1 encoding?

A: ***Coming soon.

7.2 Storage Formats

7.2.1 What is Super Audio Compact Disc (SACD)?

A: ***Coming soon.

7.2.2 What is DVD-Audio?

A: The following basic DVD-Audio description was originally printed in the latest B&H Audio/Video equipment catalog. It is quoted here without permission, and with a couple of small edits. Keep in mind that the improvements of DVD-Audio over CD are very real, but the reader may find the following to sound a little like a sales pitch. Anyway, it’s worth a read.

“The new standard for super-high fidelity sound, DVD-Audio goes far beyond the CD in virtually every aspect of audio performance, from frequency response and dynamic range, to vastly higher resolution of musical and sonic detail, and multi-channel capabilities. DVD-Audio discs look almost exactly like CD’s, but offer 7 times the amount of storage for ultra-high performance stereo sound. Music sounds live, with true definition, dynamics, and acoustical tone. It also offers multiple channel stereo sound, has versatile playback compatibility, and allows very creative possibilities for storing text, images, and video, too.

DVD-Audio discs store 4.7GB of data, or up to 400 minutes of 2-channel stereo sound at the standard CD resolution (16 bit/44.1kHz), up to 74 minutes of 6-channel sound at 24 bit/96kHz, or 2-channel sound at an incredible 24bit/192kHz resolution. Such precise reproduction means phenomenal sound quality and extreme fidelity from DVD-Audio recordings. DVD-Audio is capable of a sampling frequency 4.3 times greater than an audio CD (192kHz vs. 44.1kHz), and has 256 times finer resolution (24-bit quantization vs. 16-bit quantization. Simply put, each second of DVD-Audio can store 1,100 times the information as one second of the same recording on an audio CD (4.3x256=1,100). DVD-Audio’s combination of greater bit-depth and higher sampling frequency yields extremely high resolution and extraordinary dynamic range of up to 144dB (vs. 96dB for audio CD).”

“DVD-Audio offers spectacular 6-channel surround sound, like home theater systems. Plus, it enables each channel’s sampling rate and bit-depth to be set independently. For example, the front three left/center/right channels can be set to 96kHz with a 24-bit depth while the two surround and one subwoofer channels can be set to 48kHz with 16-bit depth. This multi-channel capability not only provides extremely realistic, spacious reproduction of ambience, it open up countless possibilities for composers and producers to better capture the acoustics of the best concert halls or cozy clubs, for reproduction right in your own living room. But because the technology of DVD-Audio is linear PCM (uncompressed data, as with audio CD's) the audio is actually much better than what you'd hear from even the best home theater systems. Instruments with higher range frequencies (cymbals, flutes, etc.) shine with a richer, more authentic timbre, and music with rapid, high-pitched passages shows excellent definition within the individual notes.

In addition to playing DVD-Audio discs, DVD-Audio players can play DVD-Video discs, video CD's, and audio CD's, so you can continue to enjoy your existing digital disc collection. Finally, the DVD-Audio format has the potential to store interactive menus and text, so you can view album and track titles, artist biographies, lyrics, liner notes, and more. User-friendly menus may let you, in the future, to choose from a variety of interactive functions."

8.0 Acknowledgements

The author(s) would like to particularly thank the participants of the laptop-tapers@yahoogroups.com mailing list for the inspiration for this FAQ, for without their continued support, patience, and obligatory input, there might be no audience. As a rapidly growing community of computer users, audiophiles, and recordists, I expect their future contributions to be pinnacle in maintaining the topicality and accuracy of the contents of the FAQ going forward.

Also on the list to thank are the kind technical support staffs at Sound Devices, Roland/Edirol, and Benchmark Media systems, who responded so timely to requests for clarification of their products.

And to all the taper section folks whose questions about 24-bit laptop recording that I didn't get to answer, I hope this FAQ shows your questions have not been ignored.

[TOP] [TABLE OF CONTENTS][Home]