stereophile

HDCD: Keith Johnson, Pflash Pflaumer, Michael Ritter

Robert Harley | Jul 10, 2018 | First Published: May 1, 1995

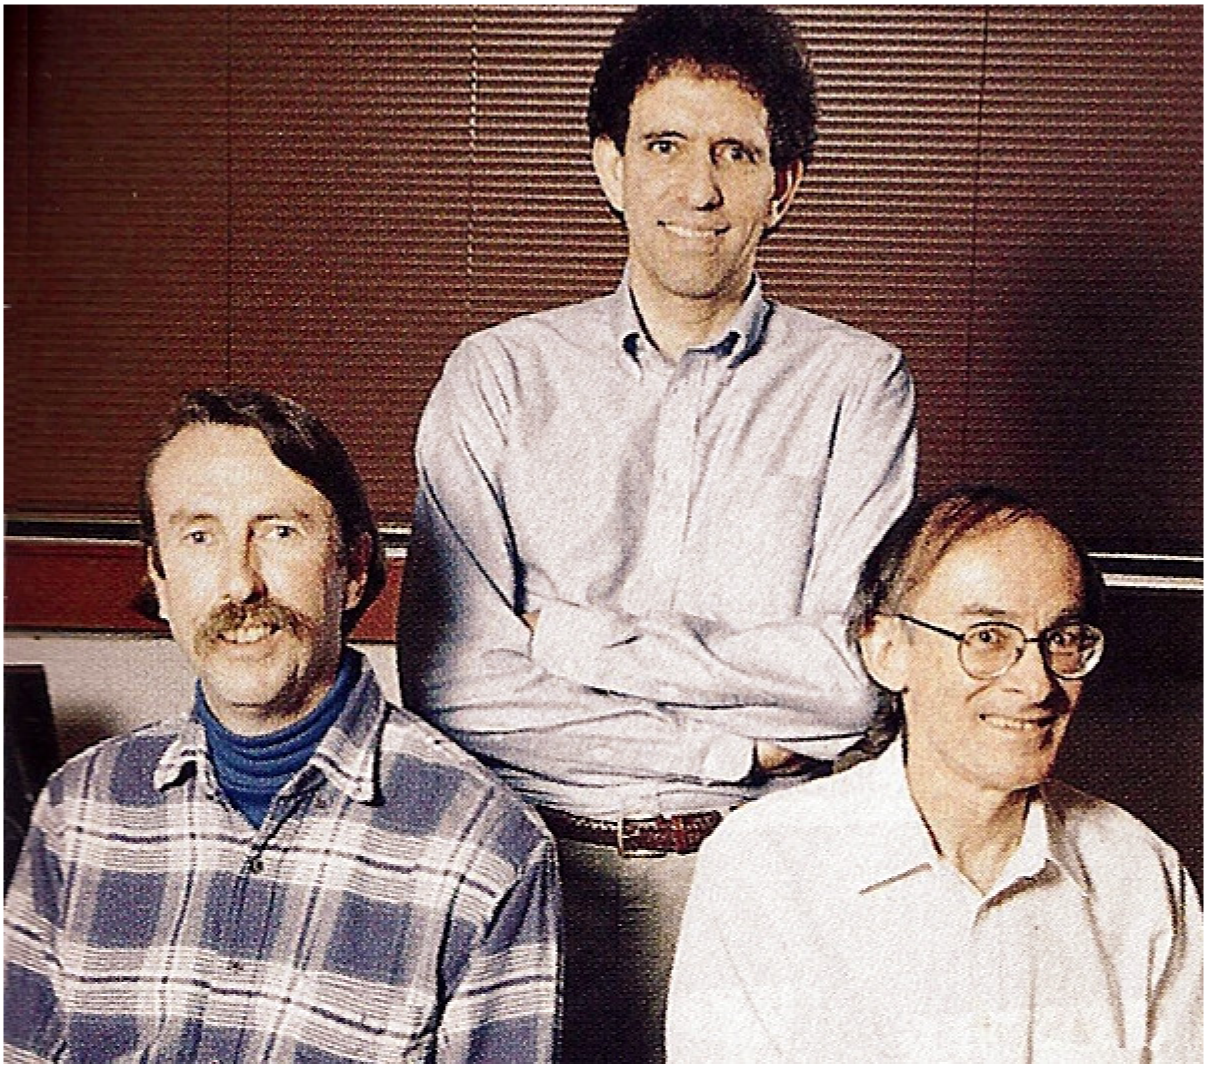

The men behind HDCD (L–R: Pflash Pflaumer, Michael Ritter, Keith Johnson

High Definition Compatible Digital® (HDCD®), the proprietary process for improving the sound of 16-bit digital audio, has finally arrived. More than a dozen digital processors using the technology are on the market, and the professional encoder used to master HDCD discs is following closely behind.

Anyone who has followed my writings on HDCD (including the review of the Spectral SDR-2000 Pro processor in this issue) knows how enthusiastic I am about the sound quality offered by this new technology. I believe HDCD is a great breakthrough in digital audio sound quality, and one that brings unprecedented resolution and musicality to home playback—all on a standard compact disc.

Pacific Microsonics, the Berkeley, California firm established to develop and market HDCD, has been reluctant to release much technical information about how the process works. Moreover, they've kept a low profile during their nearly ten-year development effort. But now that digital processors incorporating HDCD decoding are a reality, I thought it time to find out more about HDCD and the individuals behind it. I visited Pacific Microsonics at their Berkeley laboratory/office and spoke with the three principles of the company: Michael Ritter, President and business manager; and the two inventors of HDCD, Keith Johnson and Michael "Pflash" Pflaumer. I began by asking Michael Ritter for a basic overview of High Definition Compatible Digital.

Michael Ritter: HDCD is a comprehensive process for greatly increasing the fidelity of a digital audio recording. HDCD is fully compatible with the existing standard for consumer playback, the compact disc, which is fundamentally a linear format with 16-bit resolution and a sampling frequency of 44.1kHz.

We had to make HDCD compatible with this standard. Yet the goal of HDCD was to achieve a *vastly* higher level of fidelity—a level of fidelity directly comparable to the finest recording technology available—*i.e.*, first-generation analog master tape or direct-to-disc by record lathe. To do this, the HDCD process had to be a conjugate system. By that, we mean a system where all aspects of the recording and the playback decoding had to be controlled as much as possible. For that reason, the HDCD process wraps around both the A/D conversion and D/A conversion.

At the same time, as a concomitant requirement for this overall level of fidelity, we had to be able to take the HDCD process and essentially cut it in half. That is, encode the recordings, but be able to play the recordings on any standard playback equipment and simultaneously hear not only no artifacting, but a substantial improvement in fidelity over what can be obtained with a commercially available A/D converter. We've been successful in achieving that.

We have a process that, when you encode with it and play it back on standard equipment, you have a better-fidelity recording than is available through any other current digital recording method. But when you encode and then play back through equipment with the decoder, and if the playback equipment is implemented to a state-of-the-art level of performance, then you have a record/playback fidelity that is arguably as good as—or maybe even somewhat better than—any other method currently available for recording and reproducing music.

Robert Harley: Before we talk more about HDCD, let's find out about your backgrounds. Keith, I understand that you built a tape recorder from scratch when you were 14 years old.

Keith Johnson: The recorder used vacuum tubes, pieces of a Sears Roebuck Silvertone machine, and heads that I built myself. It recorded three tracks, and made master tapes that were later released (footnote 1).

When people at Ampex heard about it, I became a summer student there but still went to regular school. Later, at Stanford, I became involved with instrumentation, transistors, and things of that nature. That was going to be my specialty. I liked music a great deal, so I was also taking pipe-organ lessons.

I liked recorded sound, but didn't like the tape noise that was happening at the time. I embarked on using photolithography to make very thin materials to produce tape heads that would operate at super-high frequencies. And with tricks of focusing magnetic fields, I could get rid of some of the noise, self-erasure, and other effects. The end result was a second three-track machine that Reference Recordings still uses to make its LP releases.

Out of that work I helped form a company called Gauss Electrophysics, usually known today as simply Gauss. It still makes high-speed tape duplicators, loudspeakers, and other things. I was involved in applying the head technology to duplication of cassettes.

Harley: Wasn't that the technology which made high-speed cassette duplication possible in the first place?

Johnson: It did make it possible. Because prior to that, the frequency losses were so great that one couldn't record high-level, high-frequency signals on the tape. At that point the cassette industry was kicked off in the way we know it now. Among other things were the use of high-speed endless-loop bins, automatic loading machines, and things of that nature. Things that are now considered day-to-day methodology were all pioneered at that time.

I also liked to make recordings, and worked with the Glendale Symphony in Royce Hall at UCLA. Those were exciting times—lots of experimentation with microphones and recording electronics.

Somewhere along the way I met Tam Henderson [in 1995 President of Reference Recordings]. I was recording a small classical ensemble at a party. He liked what he saw and heard, so we did several projects and introduced a recording called *The Astounding Sound Show* [RR-7]. That essentially put Reference Recordings on the map, and paved the way for more serious work on microphones, electronics, and the recording system we have now.

Ritter: One other thing you might mention is the work you did on optical discs.

Johnson: I did a lot of work with Paul Greg, who is essentially the granddaddy of the video disc. He's the guy who had the concept of embossing plastic and reading it back with a servo—the basics of the industry. I did one key patent in that system which everybody uses today.

Harley: You also designed and built the entire recording chain for Reference Recordings, and have engineered all their releases.

Johnson: That's true.

Harley: What's your background, Pflash?

Michael "Pflash" Pflaumer: I've been interested in electronics and music all my life. In high school, I was one of a few people to put a radio station on the air. We built all of the equipment, including the console, modulator, and transmitter. We got to play our record collections on it. And I was also very interested in ham radio, microwaves, and digital encoding schemes.

Around 1960, I built a digital PCM encoder and decoder. It was only 5-bit, but it was all built out of vacuum tubes, because vacuum tubes were cheap in 1960 and transistors were very expensive. It took several racks full of stuff to make a 5-bit D/A and A/D converter and transmit the signal over a microwave link. A friend and I did this as a research project, and experimented with different modulation schemes.

Later I got involved in television production for broadcast and was the chief engineer for a studio. We did a lot of music production, where we televised string quartets and things like that. I have some exposure as a recording engineer as well as being chief engineer at the facility.

About that time I also started getting interested in computers. When I was an electronics major in college, computer science was not a separate discipline—it was all folded into the electrical engineering department. In the early '70s, when the first integrated microprocessors came out, I bought an Intel chip set and built myself a microcomputer. I basically had to write my own software and operating system.

The next big step was getting involved with computer networking. A friend of mine had started a company to make a local-area network for CPM computers. He called on me when it was apparent the people he was working with were not able to actually make the thing work. I ended up designing network hardware and writing all the network software that would run on K-Pro computers. That was called the Web. It didn't sell more than a few thousand copies.

Out of that experience came the impetus to do a local area network called Tops, which was basically the first computer network to transparently interconnect machines with different operating systems. And by the time it actually reached the market in 1986, we supported the IBM PC and the Apple Macintosh, and very shortly thereafter some workstations running Unix. And from each machine, the rest of the network, which was a distributed file server, looked as if it were part of your machine. So if you were on a Macintosh, the file system of an IBM PC was still icons and folders. It all worked transparently, including name translations from long Macintosh-style names to IBM-style names to Unix-style names, and it had a very elaborate directory translation. It turned out to be a very successful product. It's actually still needed today, unfortunately. We sold the company to Sun Microsystems (footnote 2).

During the time I was doing Tops, Keith, Mike, and I got together to try to realize what we now call HDCD. Keith had come up with the seed ideas, and the core realization that something needed to be done about digital sound quality. The standard format, even when executed properly, doesn't contain enough information to be satisfying.

Ritter: I was involved with Keith in a project at that time. There were some people interested in funding a company to produce a line of high-quality recording studio equipment. The industry was suffering from equipment with hundreds and hundreds of cheap op-amps, lots of bells and whistles, and zero sound quality. That project didn't come to fruition, but that's why I was involved with Keith. Keith mentioned some ideas he had to me which were very interesting. I told Pflash a little bit about Keith's ideas, then brought the two of them together.

Pflaumer: Keith wanted to get people interested enough to put in money to develop his ideas, but he didn't want to tell anybody what he was really thinking about. He didn't want anybody to steal his idea and do it without him, and was very careful about what he revealed. I initially came along as kind of a technical expert for some of the potential investors Mike had brought in, to see whether I thought this was real or not.

After we had our meeting with Keith, the investors asked me whether I thought it was possible. I said "Yes," and told them what I would do if I were involved. Later, we figured we could trust each other and signed non-disclosures. Keith disclosed his ideas and I disclosed mine. Because we come from somewhat different backgrounds, it turned out that our ideas fit together very well. We both thought of a couple of things in common; he thought of things I hadn't thought of, and I thought of some things he hadn't considered. When you put them all together, it amounts to what years later has become HDCD. Keith said, "Let's become co-inventors and do it." Mike said, "I think I can get the money to do it." That's how it started.

We each have a little common knowledge of the whole territory. Keith is a superb analog designer and a superb recording engineer. I have a lot of strength on the digital side of things, computer algorithms, and digital signal processing algorithms. It all fit together.

Harley: What did you find lacking—musically and technically—in conventional digital that made you think the world needed a higher-quality method of encoding digital audio?

Johnson: In a nutshell, digital didn't give you the impression of "being there." Around the time we introduced *Däfos* [initially Reference Recordings RR-12, now Rykodisc RCD-10108], digital was just getting going. Early on, we rented a Soundstream machine and were appalled at how unlistenable the playback was. It produced a grating sound that made us not want to hear the music once it was recorded. We felt it was sheer craziness, and we made an outcry about it. One of the first outcries on my part was in *Stereophile*—the first *Stereophile* that had a full-color cover, incidentally (Vol.7 No.4, August 1984] (footnote 3).

It was painfully obvious that sub-order harmonic distortion and noises were getting in. It was the result of high-frequency things creating distortion components that were not harmonically related to the lower frequencies.

If I heard that kind of thing, I thought I should be able to measure it. I devised an eight-tone cluster test, a test signal composed of high-frequency tones strategically placed in frequency such that the difference frequencies between tones are multiples of each other's and are troublesome to the system's sampling rate. I figured this would probably make a mess. Indeed, on the early converters, it did. At best, the sub-order components were maybe 60–70dB down. At that time I didn't realize how bad that kind of performance was.

I proceeded to start tearing the system apart, looking at the digital systems and discovering what was going on. Why does a thing measure so well with continuous pure tones, yet the perception of music is so poor compared to the live microphone feed?

I built my own converter based on a Sony PCM-701ES, which at the time was a pretty much state-of-the-art A/D processor. It was properly dithered, the levels were well-defined and independent of program condition, it had a very good filter—things like that. It was used to make all the earlier Reference Recordings material.

You could hear that it was adding artifacts: timbral shifts—sounds of instruments that weren't right—and the lack of smaller "environmental"-type sounds that contribute to a sense of realism. In the recording sessions we'd use the analog tape for playback because it had more going for it. More entertainment. More of the artist. More of those things that count—things that got you involved in the program. And by that time, the quantization "gritchies"—the spray-can cymbals, the cardboard bass, and other kinds of things that are wrong with digital—were pretty well shaken out of the system.

Harley: And this was even on your custom-built digital recorder?

Johnson: Yes. You'd put music into the digital machine and it was like the bone stripped of flesh: not much there. The problem wasn't just distortion, it was a lack of information as well. It was clean and had a certain degree of transparency that was somewhat better than the analog tape, but it just wouldn't bring you *into* the music. The staging would never be wide enough, never deep enough, and you couldn't sense the front-to-back space in the room. There was almost always a difficulty in identifying instrumental timbres—soprano sax, alto sax, or clarinet would have problems. It was hard to identify *how* the instrument created the sound. You heard the piano sound, but couldn't easily identify how the sound came about: a hammer striking a string. Something was missing. That was the frustrating part of it.

I then set out to try to figure out what was going on. What is it that we hear that doesn't seem to show up in measurements? And what can be done about it? You go into the system and try to find out what's wrong, and then make a measurement to see if it does in fact create a problem. If there is a measurable problem, can you hear it?

It became obvious that the dynamic range of the CD wasn't even *nearly* adequate for a good recording. With an analog tape recorder, which might have a 70dB signal/noise ratio, you can gain an extra 10–15dB by pushing the tape into overload, and you could easily hear 15–20dB into the noise. That's more than 90dB of dynamic range. But with a digital system, anytime you had more than 40–50dB of dynamics in the program, you were in trouble. The granularity and loss of information at the lower levels were intolerable.

Footnote 1: You can hear this tape machine and Keith Johnson's early work on The Red Norvo

Quintet CD *The Forward Look* (Reference RR-8CD). The recording was made on December 31, 1957, but wasn't released until 1981.—**Robert Harley**

Footnote 2: Tops was enormously popular, winning many awards and becoming the standard multiplatform networking package in the computer industry.—**Robert Harley**

Footnote 3: Actually, the very first *Stereophile* to feature a four-color cover was Vol.3 No.12, Spring 1977. Vol.7 No.4, featuring Keith, was the second four-color cover.—**John Atkinson**

HDCD: Keith Johnson, Pflash Pflaumer, Michael Ritter Page 2

You therefore had to draw a line and say you've got a 40dB-dynamic-range system; I can't use the stuff that's below that. The problem is made worse if there's a good rim shot on a drum or a cymbal crash. A peak level on an orchestral session can be 110–120dB. The analog tape will saturate, but it only happens during a brief moment in time. You're not aware that it has happened, but a digital system will go crazy if you overdrive it. After the peak, you're back in the soft passages below the 40dB range where things don't work well.

The very thing the digital systems were touted as being very good at was the very thing they didn't do. It would do signal/noise ratio, but couldn't do dynamic range.

That was my observation of the things in digital that had to be fixed.

Pflaumer: I think it comes down to the fact that human hearing is very sensitive to very small details, even in the presence of large signals. The whole area is very difficult to measure properly. Just measuring the low-level performance of a converter, for instance, does not adequately show you what the converter does in the presence of a complex signal. But the ear is very sensitive to small things, even when they're riding in the midst of a very-large-scale signal. This happens when you have all kinds of instruments playing at the same time. The details of an oboe's reed sound, for example, get lost in the conventional digital recording because the details that distinguish the reed are very small compared to everything else that's going on. In conventional conversion, even with dither, the harmonic structure tends to get lost.

This may be a good time to mention some of the limitations of dither (footnote 4). There is certainly a prevalent misunderstanding these days that all you have to do is dither a 16-bit signal and you get enough resolution to do the trick. There are a lot of people who hold that view—or the view that dither with noise shaping ought to be enough. The problem is that dither is an averaging kind of phenomenon. You only get that resolution by averaging over time. Averaging works very well for lower frequencies in the musical spectrum. The problem is, of course, that it doesn't work very well for the higher harmonics of the signal.

Dither doesn't help instantaneous high-frequency components such as percussive edges—the little peaks a reed creates, or the very sharp, time-aligned compression peaks that a brass instrument makes. The resolution is not spectrally flat, and human hearing is able to detect these things. The human ear does not average things like the spectrum analyzer does. As a result, a lot of the delicate harmonic structure that gives rise to our sense of timbre is not properly preserved. Digital audio needs more bits, basically, than 16 real bits. You can't get the same thing by dithering [a 16-bit system].

We did a lot of experimenting with both dither and noise shaping as part of the evolution of HDCD. We did a lot of listening tests and implemented various classic curves that people publish for noise shaping. And basically, we didn't like what we heard very much. There was always a sense that something was a little bit unnatural about them. The lack of ability to preserve the harmonic structure of an instrument was part of it. The other part of it, we concluded, is that, in nature, a shaped noise floor doesn't exist. If you're in a hall listening to instruments, you don't have that kind of a funny spectral shape to the noise floor. Somehow, even though you can't hear it, there is still a sense that there's something not quite natural about it—like a pressure on your head. Even though you don't actually *hear* the noise floor, there's a sense that something's wrong.

Johnson: There's another factor, too. These schemes can use a significant amount of high-frequency energy that's focused at the extreme top end of the spectrum, and if it's played back on a cheap piece of electronics, like a boombox or something with integrated-circuit amplifiers, the TIM distortion generated in these things creates havoc. Quite often, some of these schemes that push dither to very high levels defeat the very purpose of the product. Inexpensive electronics can't play it back. Most high-end systems are much more tolerant of situations like that, and have relatively benign performance.

Harley: So HDCD uses no noise shaping?

Johnson: No, we don't use noise shaping.

Ritter: There's an important point to make here, and this noise-shaping discussion is a very good illustration of it. The team that invented and developed the HDCD process is an impeccable technical team that followed rigorous scientific and technical procedures. HDCD was definitely not the sort of thing that could be conceptualized or developed by a seat-of-the-pants, play-it-by-ear kind of approach. You couldn't get to where we are now just on the basis of how things sounded. There's some very serious scientific work involved with it. However, at the same time, throughout the development, how things sounded was very much a part of the equation. Even though we did very elaborate measurement test setups, at the end of the day the bottom line was, What did it sound like?

As Pflash pointed out, if you use conventional test equipment—spectrum analyzers or FFT machines—some of these noise-shaping approaches appear to have real benefits. However, when you have a controlled listening situation and very-high-resolution source material—Keith's first-generation analog masters—you can analyze these different technical approaches. After all the research we've done, it turns out that human hearing is far more sensitive than any measurement device—even the latest test equipment we have.

We couldn't have gotten to where we are today if it wasn't for a combination of technical expertise, scientific background and approach, an extreme degree of conversance with live and reproduced sound, and that high-resolution source material. There was a synergy between these to the point where I would say that we could not have achieved what we have without *all* those elements being brought together and applied rigorously over a period of years.

Harley: Keith's experience hearing the orchestra live, then the microphone feed, then what digital did to the signal, must have been a great asset.

Ritter: Exactly. We had those references. We knew what was possible. And then we had the mental horsepower to deal with solving the problems.

Harley: When did you first start working seriously on HDCD?

Ritter: It was in the spring of '86 that Keith first described some of the concepts to me. And then as Pflash got with Keith, the synergy started. It's quite an amazing thing. Not only do their talents complement one another, but the *level* of their talents also complement too. Both of these guys are brilliant, and used to a situation where they're a couple of years ahead of everyone else on whatever they're working on.

I did the other steps of getting the capital together and forming the business. We incorporated Pacific Microsonics in November, 1986.

Pflaumer: To be fair, we were all only working on it part-time in '86 and through about '89. I was still very much involved with Tops up through '89, so I was only able to tear myself away for a couple of weekends a month to confer with Keith and try to make the HDCD ideas gel.

Johnson: At that time, we were simulating the HDCD concepts in the analog domain. We would take one part of the system, isolate it, and then build a processor or whatever was necessary to develop one little piece. When Pflash came on board full-time, it was fortuitous timing; Pflash's knowledge and experience in digital signal processing, along with the availability of powerful DSP, gave us the opportunity to take it to the next level.

Ritter: We attempted to emulate some of this stuff on the most powerful computer we could find at the time, which was a *very* expensive Sun workstation. It ended up that we needed about eight times the computing power of this RISC-based workstation to run the HDCD encoding algorithms in real time!

After '89, there was a much higher level of development activity going on. Keith was working on better and more elaborate implementations of the A/D. There was a *huge* amount to do, because everything that we were doing was just beyond what anybody had done before. There were no off-the-shelf solutions.

Pflaumer: Before we could get into a realistic, real-time implementation of HDCD as a process, we wanted to build an A/D and a D/A which satisfied us—in terms of the basic quality limitation of conventional coding with more bits and a higher sampling rate—without worrying about trying to fit it through a 16-bit, 44.1kHz pipeline. In other words, we needed to put an A/D converter and D/A converter back to back to set the quality level we were working toward.

We first built an A/D and D/A that sounded pretty good, then implemented the various concepts that had been discussed as to what HDCD should be. We then implemented those ideas in a digital form and squeezed that into a 16-bit, 44.1kHz signal we could record.

For the first time we were able to do all of the things that we thought that we should do and had previously simulated in the analog domain. We could do them all simultaneously, in real time, and were able to process real audio signals and get the kind of quality we were striving for through a conventional 16-bit, 44.1kHz channel.

Ritter: That was a pretty exciting time, because there had been *years* of slogging with nothing to listen to. When we first started getting the thing working, it was way beyond any other kind of digital. It wasn't as good as it is today by any means, but the thing was working. It was very exciting.

Harley: Specifically, what's going on in the encode process that we have been discussing?

Ritter: For a variety of reasons, we can reveal only so much. It's important to reiterate here that HDCD is a holistic system, meaning that it addresses all areas of digital recording and reproduction. It has to. If you just say, "We make things better by doing *x*, *y*, and *z*," it doesn't begin to address the overall problem that we're confronted with. Therefore, the process itself wraps around the A/D and D/A conversion and is integral to it.

In that context, HDCD begins with an extremely high-quality, proprietary A/D converter, which is arguably better than any other converter that we're aware of in any form. It's not a little better; it's a *lot* better. It's better in terms of distortion generated in the conversion process, and it also has a very high degree of resolution. It has wider dynamic range, extended frequency-domain response, more bits than 16, and a much higher sampling frequency than 44.1kHz.

The signal that we get from the A/D conversion has far too much information to record or store. This signal has all that information and very low distortion at the same time. That signal is then analyzed using DSP techniques in real time. The algorithms that look at the signal are algorithms that were derived from our research into psychoacoustics and auditory physiology. We were concerned with how we hear mechanically in addition to how we hear subjectively. Those algorithms look at this high-definition signal and determine the components in the signal that would not fit in the normal 16-bit, 44.1kHz recording—signal components that are important in terms of how we perceive subjectively and also how we perceive objectively. The high-resolution signal is then decimated to a 16-bit, 44.1kHz signal that can be recorded on a compact disc, but with the additional information about the psychoacoustically important signal components added in.

The additional information is added in two fashions. Part of it goes into the linear PCM signal itself in a way that can be reproduced to a certain extent with standard playback equipment.

You can hear some of this improvement and some of this additional information on standard playback equipment. Second, additional information goes into a buried control channel in the LSB [the 16th and least significant bit of each 16-bit audio sample]. The buried control channel doesn't occupy the entire LSB; it's done in a very clever fashion, occupying only a very small percentage of it. This is Pflash's work here.

Pflaumer: The encrypted control channel shares the LSB with the LSB of the music. One of the key ideas here is that the additional information is not needed in a steady-state fashion. There are certain times in the program material when you need to provide much more information than at other times. The side channel can share the LSB with program material, and gets inserted as needed. The decoder is watching for its presence and picks out commands as they're sent across through the channel.

Ritter: It's quite amazing—almost like something for nothing. You literally have this additional information sufficient to reconstruct the original high-resolution signal. However, we don't do it by taking away any resolution in the non-decoded playback. Essentially, there is no loss. On average, the additional information uses only one to five percent of one bit.

Harley: But is that a high enough data rate to transmit the reconstruction information?

Johnson: Yes, it is. In the encoder, we determine the process which works best, then send the information of which process was used down the control channel. On the playback side, the decoder says, "Ah, that's the process I need to perform to be the conjugate to the encode process."

All it needs to be is a number that says, "This is the process to perform." In the meantime you've got a powerhouse—the decoder—at the other end that's not creating information, but is programmed to do some fairly complicated activity.

Pflaumer: You can send with brute force lots more information in that channel, which may be desirable at certain times. But for the most part, because HDCD is a process that has anticipated the requirements under different conditions, the encoder can pick the appropriate process based on the analysis of the signal, and simultaneously tell the decoder which process it's picking. The decoder knows how to complement that operation. As a result, HDCD provides the equivalent of a lot more information without having to have the bandwidth.

Harley: You transmit the *command* to perform the restorative operation rather than the operation itself.

Johnson: Exactly. You could, for example, send a code that says, "Output Beethoven's Ninth." Because the receiving end has Beethoven's Ninth in memory, I can play back Beethoven's Ninth with just a few bits. It's an extreme case, but it makes a point about how HDCD works. It's a powerful technique that you could never do in the old analog system.

Harley: Does the presence of this encrypted channel degrade the fidelity when played undecoded?

Ritter: No. That's why HDCD discs sound so good on standard playback. We're not throwing away anything.

Harley: What types of musical information that would get lost in the 44.1kHz, 16-bit bottleneck do you encode in the control channel and then reconstruct on playback?

Ritter: Timbral information, hall ambience, low-level information that gives you accurate timbral reproduction of instruments and voices. The additional information also preserves spatial cues.

Johnson: A lot of the things that you can't get by dithering. We developed other ways than dithering to preserve those things. A lot of what we do is looking at the signal on an instantaneous basis—what a spectrum analyzer *doesn't* do.

Ritter: That's an important point. HDCD preserves instantaneous information in the signal, in terms of what's perceived as frequency extension.

Harley: By perceived frequency extension, I understand that to mean that HDCD can preserve components in the signal that trigger a response in human hearing equivalent to a frequency response extending beyond 20kHz.

Footnote 4: Dither is a small amount of noise—either broadband white noise or narrow-band noise—added to the signal to randomize quantization error, allow the digital system to resolve information lower in amplitude than the least significant bit, and make digital encoding more linear.—**Robert Harley**

HDCD: Keith Johnson, Pflash Pflaumer, Michael Ritter Page 3

Pflaumer: That's true. In the human ear, the cochlea has hair cells that respond to about 80kHz. And yet, for steady-state sinewaves, even a young person can hear only 20kHz as a steady-state tone. As we get older, that deteriorates. But the sense of frequency extension is largely based on transient events—things that happen only in short instants of time.

The shape of a wavefront generates harmonics in the bone structure of the middle ear, which are then picked up by the hair cells in the cochlea and used to reconstruct the picture of that wavefront.

The army did quite a lot of research on perception of this kind of supersonic wavefront. It has a lot to do with spatial perception. The army wanted to understand how the location of gunshots can be pinpointed by listening. If you disturb the shape of the wavefront going into the ear, you very greatly disturb the ability to pinpoint the source of the sound. A lot of what we do is pay attention to these transients, which are very brief instants of time as opposed to steady-state high-frequency response. It turns out that the transient information is much more important than frequency extension.

Harley: I've noticed with HDCD an ultra-fine resolution of image placement in the soundstage.

Ritter: Yes, that's exactly right.

Harley: Can you describe in more detail how this transient information is reconstructed on playback?

Ritter: You can see why we're sort of closed-mouthed about these things. People don't know about this stuff. It's just one of the things that makes HDCD work. We don't want to give away some of this information.

Pflaumer: After the professional encoder is out there, if someone seriously wants to reverse-engineer it, they probably will. At that point, it's appropriate for us to do a full-blown AES technical paper on everything that's going on (footnote 5). But we're trying not to give the competition any head start at this point. We worked real hard for this information.

Harley: Going back to the hidden control code: the fact that you use the LSB only briefly avoids the problem of the control code being correlated with the music, which would tend to make it audible when played back undecoded.

Pflaumer: The encrypted channel uses spread-spectrum–like techniques. Even though the information that's in there might have distinct patterns, it's randomized. By the time the encrypted channel is actually inserted in the LSB, it has as close to random characteristics as you can get.

Harley: Is there any chance that the encrypted data could create problems for some D/A converters?

Pflaumer: No.

Harley: What's the availability of the HDCD decoder chip? If I were a digital-processor manufacturer and wanted a thousand pieces, when could you deliver?

Ritter: We are in full commercial production. You could get a thousand pieces in mid-February 1995. We have more than 17 licensees.

Harley: Tell me more about the algorithms that examine the high-resolution digital signal to determine what information to put in the encrypted control channel.

Pflaumer: Keith and I together formulated the original algorithms in very broad terms. As Keith mentioned, some of them were implemented as analog models—pieces of the algorithm, very early on, which Keith did. I'm responsible for taking all of those and turning them into digital algorithms and writing the DSP code to implement them. I've been working on this full-time since the beginning of 1990.

All of the DSP code is done in assembly language because it's all very time-critical. The prototype encoder had 11 Motorola 56001 DSP chips, and the algorithm was shoehorned into those 11 processors on the encode side. Everything was done in assembly code because the higher-level DSP compilers are too wasteful of cycles. The professional encoder will use the new Motorola 56007 processor running at 66MHz.

I also did the algorithm that is incorporated in the decoder chip, the PMD100. We had a consultant do the actual schematic design of the chip, but I did a mathematical specification, bit for bit, of what the chip should do in mathematical terms. Then I implemented that on the DSP engine, and we used that to generate the production test vectors for the silicon, and for the simulations before the silicon. The DSP implementation allowed us to hear playback precisely, bit for bit, through our part, long before we even had schematic diagrams of the chip, much less actual silicon. So we knew what it was going to sound like long before the chip actually existed. Nothing was left to chance.

Harley: Is it possible to make improvements in the encoding algorithm, yet keep the same decoder? If you get a large base of installed decoders, will they benefit from future improvements, or are the encoding algorithms fixed?

Pflaumer: Anything is possible. We are at this point pretty satisfied with the encoding. As we make improvements in our own D/A conversion, we hear more and more things in recordings that we've already made.

It's somewhat analogous to the situation with LPs. There are many great LPs, but nobody could actually hear everything that was there when they were made. It's only at the tail end of the life of LP playback equipment that it has reached the limit. People are still surprised at how much information is there.

This is not the case with most conventional CDs. The size of the information stream is so limiting that the better of the CD playback equipment reaches a certain threshold of performance. When you get to a perfect conversion of 16-bit CDs, then there isn't anything more you can do. We put in enough additional information that the D/A converters have not quite caught up yet.

Ritter: We're challenging conversion technology, particularly on the D/A side. As good as HDCD discs sound decoded today, they're going to sound even better three or four years from now.

Pflaumer: We're definitely upping the ante for all of high-end audio. Now there's a signal source with a resolution that requires doing everything right in order to hear that resolution. It's amazing how much is there. It surprised *us*.

Harley: The professional encoder used to make HDCD-encoded recordings is a critical element in the commercial acceptance of HDCD. When will it be available?

Ritter: We're looking to begin shipping pro encoders in June. We'll have parts of it working before that. Some of the earlier versions of it may go to a few very interested parties who are champing at the bit, and who are in very powerful positions in the industry.

The pro encoder is the entire thrust of the development effort at Pacific Microsonics today. We hired a significant full-time staff. We brought in René Jaeger, who has spent much of his career bringing state-of-the-art digital audio products to the recording industry [at Lexicon and other firms].

The level of accuracy and of distortion that the professional encoder is designed to achieve has never been done before in a digital audio recording product. We even developed our own discrete hybrid op-amp modules after investigating *everything* that was commercially available. We looked at the most expensive commercial monolithic op-amp parts and found them wanting.

We're not compromising on this as a product. It's going to be a product that will set the standard for—I was going to say digital audio recording. I'm going to say *any* audio recording. When this becomes available, I think it's going to set a new standard for the kind of quality that can be achieved in studio recording. To achieve that level of performance is quite a project. But we have the best possible people working on it, and we're quite far along in the development of it.

We have a number of encoders already sold to some of the top facilities in the recording industry. One of the largest mastering houses in the country, Georgetown Masters in Nashville, has ordered a half a dozen of them. A number of other people that I probably can't mention by name have ordered encoders—and these are very highly placed people in the recording industry. The fact that we're having such a high level of appreciation and establishment at that level bodes very well. This product has so much quality to it that once we get units to people that are doing classical recordings—and not just the esoteric labels, but the name labels—and they hear what this can do, we're quite confident that it's going to be recognized as a breakthrough. The more challenging the material, the more apparent the superiority.

Harley: What are the prospects for HDCD encoding on major labels?

Pflaumer: One of the surprising things to a lot of the recording engineers who have heard HDCD in a real studio environment is how neutral we are. They spent a lot of time getting a particular sound in a mix. When you take that final two-channel mix and record it on our encoder, then play it back, you get back what you put in. This is a tremendous surprise, because it's not true with any other A/D converter currently in use in the pro industry.

We were at a major recording studio in Los Angeles—I probably shouldn't say who—and other engineers who weren't involved in our project kept wandering in and listening. In many cases they couldn't tell which was the HDCD playback and which was the feed out of the board from the multitrack machines. You couldn't tell which was which. They'd never experienced that before with any kind of recorder.

We did some simultaneous tests with 30ips analog recordings. The analog sounded very, very good, but it had a character to it which told you it was an analog recording after final mix. And we did the same thing with good-quality studio A/D and D/A conversion. I don't want to mention names here, but the difference was dramatic.

People hear it and say, "When do I get mine?" "How do I get one?" "Am I in line?"

Harley: Do you see mainstream releases being HDCD-encoded from mastering engineers making the process available to individual artists, or record companies adopting HDCD wholesale?

Ritter: I think both mechanisms are going to occur. There's a third mechanism as well. There are a number of artists who will insist on HDCD before they get to the mastering facility. They'll insist upon it not necessarily because they're audiophiles, but because they want to preserve the quality of the music they've spent hundreds of hours creating. Listening to the output of the multitrack machine from the console, they get a presence, a vividness, and a palpability to the sound that are destroyed when they make the two-track digital master. The life and presence are taken out of the recording; the impact, the emotion, is gone.

But the artists and producers are very tuned-in to those things. They know that it motivates people to buy recordings. They see HDCD not only from an aesthetic point of view, but from a commercial point of view. Mastering facilities will demonstrate and recommend HDCD, artists are going to demand it, and it's possible that record companies will adopt it.

Johnson: One performer brought his master tapes here that had guitar solos. You think, okay, a guitar solo is a guitar solo. The guy is wailing into the amplifier, and the thing is about to blow its gourd from distortion. You think that a digital system isn't going to mess that up because it's already so distorted. How can anything more happen to it?

But in the hands of a very fine artist, that distortion has a real meaning. You can get into it. There's a very complicated tapestry of sound going on, and the guy's a master at weaving and controlling it. When that sound goes through a digital process that isn't right, the distortion is so horrible that you can't stand to be in the room. The musical intent is lost. It's not just the intent—there's nothing there.

I was shocked the first time I was exposed to this. I always thought that the studio people were going to have some problems using something as good as our encoder because a lot of their equipment has problems. It's been like using a Brownie camera. You spent $50,000 making a recording, and then you pass it through this little $24.95 A/D part. You can crank through hundreds of rolls of film, as it were, and 50 years from now that film is not going to be worth very much, as opposed to something that does it right and comes back the way it went in.

Pflaumer: For the first time, we let them get back in the final, distributed playback medium exactly what they put in. They don't have to mix by listening through the imperfect chain and trying to compensate for what's wrong with it.

Ritter: It's been wonderful for us to use Reference Recordings releases, but what you're hearing is not only HDCD, but Keith's recording technique and his whole approach. When you start getting commercial HDCD discs, they're not going to sound like Reference Recordings. Some of them will sound excellent, but they'll all sound different.

For example, some of the transfers we've already done using old analog master tapes are very simple recordings. You could almost say they were crude. But when you listen to the analog master, there's a sense of people in a room doing something. And that's preserved through HDCD.

When you went through a conventional digital transfer, the warts remained and the musical stuff got trashed. All you really had was a compendium of digital artifacts and the warts and errors of the recording. It's a very interesting thing to hear.

The point is that HDCD is neutral. HDCD recordings don't sound like Keith Johnson recordings, as wonderful as those are. They sound like whatever you put into it. You're going to start seeing a whole spectrum of different kinds of sounds coming out of HDCD once the encoders get out there.

Harley: A limitation, though, is that any signal that has ever been digitized with 16-bit linear encoding won't benefit from HDCD.

Ritter: Yes. If you have a 20-bit master, we can preserve the dynamic-range information. But you're stuck with the distortion and the frequency-extension limitations.

Johnson: Then there's not much you can do. The information is already lost.

Harley: Are you worried about the "not invented here" syndrome creating resistance to HDCD in the industry?

Ritter: Even though Dolby noise reduction wasn't invented by any large electronics company, it became ubiquitous in the marketplace because it works. Because people who buy the products demanded it, the electronics companies really had no choice.

And Ray Dolby was reasonable in how he did business with those companies. He gave them a situation where he got a small royalty out of a large market, and did quite well. It wasn't painful for them to do business with Dolby—and we'll be the same way.

We think that any company, even a very large company, that takes a clear-eyed look at the bottom-line impact of using HDCD technology will recognize that their products will sound better and they'll make money from it. For all those reasons, we think HDCD could migrate essentially to every area of the market. In fact, there's no reason why it can't become a standard.

Footnote 5: See http://www.aes.org/e-lib/browse.cfm?elib=7387 —**John Atkinson**

HDCD: Keith Johnson, Pflash Pflaumer, Michael Ritter Page 4

Johnson: Because recordings made with the encoder sound better on standard playback equipment, there's no hindrance to the recording industry for using HDCD, even though the decoders are not out in large numbers yet. And since HDCD-equipped playback equipment makes standard CDs sound better, there's a real advantage to putting the decoder chip in playback equipment, even though there aren't a lot of HDCD-encoded recordings out yet.

The combination of those two things, combined with the magic that's there when the encoded recordings are decoded, will allow HDCD to enter the market from both ends. HDCD can move into the industry easily with nothing but benefits for everyone involved.

Harley: Many digital audio workstations used to prepare CD master tapes change the data as a matter of course, which would destroy the encrypted control channel and prevent HDCD decoding. Is this a significant concern?

Ritter: Yes. But luckily, all the available equipment has modes that pass the data uncorrupted. If you start throwing dither on the LSB, and that sort of thing, it's goodbye to the control code. [*When I compiled the master for* Stereophile*'s Test CD 3, which has one HDCD track, I was very careful to do all transfers in the hard-disk editor's bit-for-bit mode, thus preserving the integrity of the LSBs.—John Atkinson.*] If the HDCD light comes on [on a digital processor's front panel], you know you have data integrity.

As we increase our marketing and sales effort, we'll be talking to the manufacturers of editing workstations to stress that the data can't be corrupted.

Harley: How do jitter-reduction boxes, DSP-based loudspeaker and room-correction systems, and other consumer data-manipulation devices affect HDCD decoding?

Pflaumer: Some of the dejittering boxes that use interpolation filters internally will end up destroying the ability to fully recover the HDCD.

I should say something about the tendency to try and make comparisons between HDCD decoded and HDCD "not decoded" by somehow messing up the thing so the HDCD light doesn't come on. These tests are not really valid because not all of the decoding depends on the presence of that buried channel. There's no way to completely shut off the decoding in the HDCD decoder chip. Even if the light doesn't light, there are still complementary processes going on without the control code.

The only true test of HDCD *vs* non-HDCD is to compare playback of an HDCD-encoded disc through a D/A converter that exists in two forms: one with our decoder chip in it and one with a standard digital filter.

Ritter: When you're running HDCD-encoded information through these devices, it's best if they leave *everything* alone. If they're reclocking in a way that doesn't alter the data, that's fine.

Pflaumer: That's exactly right. If you want to reproduce HDCD to the highest level of fidelity, don't monkey with the signal. If you want to stabilize the time base, that's fine, but don't alter the data in any way.

Harley: One thing that surprised me about HDCD-based processors was how much better they sounded than NPC-filter–based processors on conventionally coded discs. What's going on in the filter section of the HDCD decoder chip that makes such an improvement on standard CDs?

Pflaumer: There are a number of factors, but we can't talk about all of them just yet. We can say that the calculations are performed to a higher precision than they are in most other filters. The coefficients are actually equivalent to 27-bit accuracy. It turns out that those last few bits, which add components that are on the order of 115–120dB down, create differences that are clearly audible.

Ritter: We have excellent stop-band rejection—better than 120dB. Another thing that the HDCD chip has which the NPC and other monolithic chips don't is a wide selection of dither on the output of the chip. There are eight selectable dithers that are designed to overcome the quantization problems of going through the filter, and to optimize the performance of the DAC.

If you inject fairly large amounts of dither, it randomizes the DAC nonlinearities. Instead of having DAC nonlinearities which are correlated to the music waveform, the DAC nonlinearities become random.

The dither we use is all above 22kHz. Since you have an oversampled signal at the DAC, you can put the dither energy outside the audioband, where it gets filtered out by the analog low-pass filter which follows the DAC. Using dither like this dramatically improves the sound from multibit DACs.

We also provide variable timing of the de-glitch signal, which can be used by some DACs.

Harley: It seems with all these variables that the HDCD decoder would be more implementation-sensitive than an NPC filter—you'll get a wider range of sonic quality depending on how well the designer has used your chip.

Ritter: I wouldn't say "implementation-sensitive." I would just say there's more potential for the competent designer to achieve higher levels of performance. In the default mode, our chip is just as easy to use as the NPC.

Harley: Pacific Microsonics certifies every processor design that uses the HDCD decoder chip. How does that certification process work?

Ritter: As part of the licensing agreement, the manufacturer must submit a production sample of a product—or a pre-production sample that's identical in every way to the production unit—to Pacific Microsonics for approval before the product is issued for sale.

We look at the fundamental implementation of HDCD decoding, and that the indicator light and trademark are used properly. And we also look at a variety of performance parameters on an advisory basis. We do not grade units, if you will. We might have our own opinions, but we keep them very much to ourselves. If the HDCD decoding is done properly and there are no other gross performance anomalies, then the unit is approved for sale. Of course, there's still a wide range of performance that's going to be achieved with various products.

We made a decision very early on that we weren't going to become a design service. The companies that are leaders in the industry are leaders because of hard work and investment on their parts. That's their competitive edge.

Harley: I know you developed some proprietary test signals and measurement techniques for digital systems. How close are these measurements now to predicting sound quality?

Johnson: There are still things that we hear that simply don't show up in measurements—even 120dB down with complex signals. We also have tests that reveal the conversion accuracy through clocking, which is different from looking at the clock and then assuming whatever comes out the converter is going to be timed right. In other words, we can look at an analog signal and identify how much jitter has been involved in the construction of that signal. And there are some surprises.

Pflaumer: Just having a good clock is not sufficient to end up having a properly clocked conversion. You can't do a properly clocked conversion without a proper clock. But there are a lot of other things that can go wrong. We don't want to disclose the test right now.

Harley: How significant a factor is jitter in the musical quality of digital audio?

Johnson: I've always looked at it like 10V *vs* 10µV, or 120dB. If you take the leading edge of a 20kHz signal and move it, how much do you have to move that waveform to create 10µV of error? It's not very much—a couple of picoseconds.

Harley: Jitter in the single-digit picosecond range is audible?

Johnson: Yes, but not necessarily. You don't have to have that kind of timing accuracy at all times and for all frequencies. The worst possible case is where the waveform is the steepest one, and where the timing event has just happened. A few picoseconds is as much error as you can tolerate to prevent hearing distortion components that aren't masked. If the jitter is higher than that, you may have the equivalent of 100% distortion because you don't hear the signal that created the distortion; there's nothing covering it up.

Pflaumer: Because HDCD has more resolution than 16 bits, reproducing the signal accurately requires proportionally lower jitter. We've gone to real extremes to be able to reconstruct clocks accurately. And it's absolutely essential. I'm glad to say that jitter is becoming widely accepted as an important factor. It's something that we've been paying attention to for many years, because if you don't pay attention to it, you cannot come anywhere close to the kind of quality that HDCD is capable of.

Harley: Did you realize the goal of HDCD?

Pflaumer: It actually surprised us that we achieved higher resolution and more transparency in the HDCD process than we were shooting for. We didn't realize what we had until we did a series of live recordings with real mike feeds, because there isn't any other signal source that's like a mike feed—especially Keith's mike feed. We did most of our research with outtakes of Keith's first-generation analog master tapes. We couldn't have gotten where we did without those outtakes.

But when it came right down to it, Keith's live mike feed had so much more information that it truly surprised us how much of that information made it into the recording. As we improve our D/A converters, it still surprises us how much more like a mike feed HDCD sounds (footnote 6).

Footnote 6: In 2000, the HDCD technology was purchased by Microsoft—see Wikipedia. Versions 9 and above of Microsoft's Windows Media Player app are capable of decoding HDCD.—**John Atkinson**