

John Chowning

Part I Of An Exclusive Interview With The Father Of Digital FM Synthesis. By Tom Darter.

From Aftertouch Magazine, Volume 1, No. 2.
Scanned and converted to HTML by Dave Benson.

AS DIRECTOR OF THE Center for Computer Research and Musical Acoustics at Stanford University [CCRMA], John Chowning has long been an articulate and enthusiastic spokesman for music produced by electronic means. While still a graduate student in composition at Stanford in 1964, he became interested in electronic music. Since the school had no analog synthesis equipment but did have a large computer, he jumped directly into digital synthesis.



In the '70s, the results of his research in the field of digital FM synthesis were licensed to Yamaha, and the rest is history. As you will discover from this exclusive, two-part interview, Dr. Chowning's keen insights are important for any student of FM. In part 1, Chowning explains his early work with FM, and describes the convoluted pathway that connected his original theory of digital FM to its current pinnacle of commercial success - the Yamaha DX7 synthesizer.

TD: When did you first start working with digital FM [frequency modulation] seriously?

JC: I guess there were two little chunks. The first step was in 1966-67, when I explored FM in the essential algorithm forms - parallel modulators, parallel carriers, and cascade. At the same time I was also working on spatial illusions. That was coming to a point of usability, so for a couple of years I concentrated on that, and didn't do much with FM.

Then in about 1970 I remembered some work that Jean-Claude Risset had done at Bell Labs, using a computer to analyze and resynthesize trumpet tones. One of the things that he realized in that work is that there is a definite correlation between the growth of intensity during the attack portion of a brass tone and the growth of the bandwidth of the signal. For the first few milliseconds, what energy there is mostly around the fundamental; and quickly, as the intensity grows during the next 30 or 40 milliseconds, more and more harmonics appear at a successively higher volume. I thought about that, and I realized that I could do something similar with simple FM, just by using the intensity envelope as a modulation index.

That was the moment when I realized that the technique was really of some consequence, because with just two oscillators I was able to produce tones that had a richness and quality about them that was attractive to the ear - sounds which by other means were quite complicated to create. For example, Jean-Claude had to use 16 or 17 oscillators to create a similar effect using additive synthesis techniques.

TD: What was your next step?

JC: At that point, I really hooked into it and wrote the first paper. Max Mathews [of Bell Labs] was astonished that such rich tones could be synthesized so easily, and advised me on a rewrite to make the presentation a little more compact. The result was finally published in 1973 by the AES [Audio Engineering Society]. It has also been reprinted a couple of times in Computer Music Journal. That paper really settled it in my mind. After writing it, I felt that I really understood what was going on - how to predict it, how to use it in some ways.

TD: How did Yamaha become involved?

JC: Right after I crystallized the ideas in my mind, the Office of Technology Licensing at Stanford [OTL] thought there might be some commercial interest in the idea. The University contacted the obvious people: The major organ manufacturers in this country, and perhaps some synthesizer people, too. Representatives from a few manufacturers came out. One of the American organ companies expressed a fair amount of interest, and sent engineers a couple of times. Ultimately, their engineers decided that it just wasn't practical for them. Frankly, I don't think their engineers understood it - they were into analog technology, and had no idea what I was talking about.

Then the Office of Technology Licensing put a graduate student from the Business School on the case. He did a little research, and discovered that the largest manufacturer of musical instruments in the world was Yamaha. This was a bit of a surprise, since they did not then have a significant market in the U.S. at the time. The officer at Stanford's OTL wrote to Nippon Gakki in Hamamatsu, Japan. It happened they had one of their chief engineers visiting their American branch at the time [Yamaha International Corporation, or YIC, located in Buena Park, CA], so he came up to Stanford for the day. In ten minutes he understood; he knew exactly what I was talking about. I guess Yamaha had already been working in the digital domain, so he knew exactly what I was saying.

TD: When did Yamaha actually take out the license on your concept of digital FM?

JC: I think it was around 1975 when they actually took the license, but they started working with us before that. They put a few good people working on it right away, and the first result was the GS1, which was of course much more expensive than they wished it had been.

TD: Why was that?

JC: It was filled with IC [integrated circuit] chips. I think what was going on is that while they were developing ideas about the implementation of digital FM, they were also - quite independently - developing their own capability to manufacture VLSI [Very-Large-Scale Integration] chips. It was the convergence of these two independent projects that resulted in the first practical instrument, which was the DX7. The GS1 was probably one generation of chip technology older, so they had to use many more chips than they ended up using in the DX7 - something like 50 to 2. Of course, that's not a one-to-one correspondence in power, but it's not too far off. I think Yamaha deserves a whole lot of credit for getting that VLSI implementation going in such an effective way.

They also added some very clever things in their implementation of the algorithms; things that were not obvious - not quite straightforward in the way one would usually work on a computer - in order to gain efficiency and speed. The consequence is that the bandwidth of the DX7 gives a really brilliant kind of sound. I guess there's something like a 57kHz sampling rate in the DAC [digital-to-analog converter]. The result is far better than we can get with equivalent density on our digital synthesizer here at Stanford. When we are running 96 oscillators, which is what the DX7 has, we have a maximum sampling rate of around 25kHz to 30kHz. That's only about 12kHz or 13kHz effective bandwidth. The DX7 is better than that, and I think it's quite noticeable.

TD: So Yamaha created the first hardware implementations of your basic idea?

JC: That's right. As far as I know, there were no digital hardware devices realized before Yamaha's first prototypes.

TD: In your early digital FM experiments, back in 1966-67, was it impossible to use envelopes to determine modulation indexes?

JC: No. I was using envelopes in that way, but I hadn't made the conceptual breakthrough. We had done bell tones: As the intensity falls away, so does the modulation index, so you go from complex inharmonic tones to essentially a damped sine at the threshold. The breakthrough for me was the realization that there is always such a strong coupling of bandwidth and intensity in most sounds, and that it is extraordinarily easy to implement that effect using digital FM synthesis.

TD: Did you develop a large vocabulary of digital FM algorithms in the early stages of your work?

JC: No, I was working with the basic forms: Simple FM, involving an FM pair of one modulator and one carrier; parallel carriers, where one modulator branches off into several carriers; parallel modulators, where a number of modulators feed into one carrier; and cascade, where a number of modulators are stacked, and you have modulators modulating other modulators. These were the basic things I had tried. I realized that the idea was definitely extensible - not just some uniquely useful synthesis technique which would do bell sounds and nothing else. I was quite sure that it was extendable. The fact that you could alter the algorithm in all these different ways and have different kinds of power was obvious to me.

TD: Did you define the basic FM algorithms for the DX7 and DX9?

JC: That was pretty much their decision. I talked to them a lot in the early days about the importance of things like key scaling, which is I think fundamental to the power of the DX7: With one function you can change the bandwidth as you go up, just by scaling the function of pitch. This is particularly important in the digital realm, I knew, because of aliasing: When you're trying to produce frequencies that are at or above the half-sampling rate, they can reflect down in unfortunate ways. With key scaling, it is easy to reduce the signal to essentially a sinusoid at the highest notes, which is ideal for dealing with the problem of aliasing. As for the design of the algorithms, some of the choices on the DX7 were really surprising to me. They all do use the four basic forms I worked with, but the implementation is different. On our system at Stanford we would have done it a slightly different way. We would have coupled two simpler units together to achieve more or less the same thing. And they've done it all in terms of six operators.

TD: Are six operators enough?

JC: Many people have asked me that question. I think it will be a long time before the possibilities of six operators are exhausted. Given the additional complexity of adding more operators, it becomes harder and harder to envision the acoustic result. Six may not be quite enough, but it is certainly richer terrain than anyone's going to get tired of in the short term.

TD: The interesting thing about the six operator question is that, unless you want more than four operators cascaded in a stack, you can get more than six operators by having two instruments and programming them as one.

JC: True. And frankly, I don't think there's much use in more than four in a stack. It's awfully hard to envision. With a linear increase in the number of operators, there is a geometric increase in spectral complexity. So going from one carrier and one modulator to one carrier and two modulators, it's not just twice as complex - it can be many times more complex. And going from two to three modulators, it almost becomes factorial. With three modulators in a cascade, there is an incredible increase in timbral complexity: With a significant amount of output from any of those operators in the cascade, you can very quickly

approach some sort of noise, because the density of the spectrum becomes so great. It also depends a bit on the ratios of frequencies.

TD: What I'm hearing, which is a very important thing for you to say, is that people don't yet realize what it is they've got under their control with six operators. Yet here they are asking for more operators - the numbers game, like having more voices in memory. They haven't stopped to look at what they already have.

JC: I think the key to understanding the instrument involves a whole lot of work on basic controls like key scaling, modulation, and velocity sensitivity, all with a simple FM pair - just two operators. Once a simple FM pair is understood, it's a lot easier to make use of the various combinations.

Part 2

Of An Exclusive Interview With The Father Of Digital FM Synthesis. By Tom Darter.

From *Aftertouch Magazine*, Volume 1, No. 3.
Scanned by Dave Benson.

AS DIRECTOR OF the Center for Computer Research and Musical Acoustics at Stanford University [CCRMA], John Chowning has long been an articulate and enthusiastic spokesman for music produced by electronic means. While still a graduate student in composition at Stanford in 1964, he became interested in electronic music. Since the school had no analog synthesis equipment but did have a large computer, he jumped directly into digital synthesis.

In the '70s, the results of his research in the field of FM digital synthesis were licensed to Yamaha, and the rest is history. As you will discover from this exclusive, two-part interview, Dr. Chowning's keen insights are important for any student of FM.

In part I [published in last month's issue of *AFTERTOUCH*], Chowning explained his early work with FM, and described the convoluted pathway that connected his original theory of digital FM to its current pinnacle of commercial success. In this month's installment [part 2], Dr. Chowning discusses his recent work with the DX7, and outlines ways in which the DX instruments can be used as teaching tools in the fields of acoustics and psychoacoustics.

TD: When did you first work with the DX7?

JC: You mean the first time I actually sat down and worked with it for more than a few random hours?

TD: Yes.

JC: In January of 1985, at IRCAM in Paris. I had seen prototypes in Japan, but that was more listening, talking about what was being done, and giving suggestions, rather than sitting down and working with it myself. And in fact they were at that time far away from what the DX7 has become.

TD: So for all this time since the DX7 has come out, your work has continued to be on the large mainframe computer at Stanford?

JC: That's right. This year marks the first chance I've really had to sit down and get to know the DX7 rather intimately. I'm writing a piece for two virtuoso pianists, each playing a KX88 controlling a TX816, with computer control between voices. Ever since the GS1, which is much less flexible than the later instruments, I felt that the technology at least coupled basic musical gestures in effective musical ways - velocity did something now beyond just making it louder; it affected the bandwidth, the spectrum. I felt that the

instruments were ready for at least a few pieces to be done with them. Now with the TX816, I think there are a lot of pieces that could be done with MIDI control keyboards. It is a different but complementary medium to that which we use here at Stanford. I plan to finish the piece soon.

TD: Is your SLAPCONGAS patch [presented in the October '85 issue of AFTERTOUCHE] the very first DX7 sound you ever came up with?

JC: Yes.

TD: Were you aiming specifically at a drum patch when you began work on the sound?

JC: Yes. That was rather purposeful and successful. I was interested in that because of David Wessel. He is a kind of musician/scientist/mathematician at IRCAM, and he is also a drummer. In fact, he was my student when he was doing his graduate work here at Stanford. We were talking about flams in the context of synthesis, and I thought it would be possible to do that with the DX7, because you have independent sections in many of the algorithms.

So I tried to make flams, I think, and then got into the idea of doing a conga, which sometimes uses flams for fortissimo sounds. I did it in about a day and a half - I'm still not nearly as skilled as some of those who've been working with the DX7 for years, but I guess my theoretical understanding was quite a bit of help. So I sat down and put together a drum which is conga-like. On all drums, pianissimo sounds are very different than a whacked fortissimo, and I wanted to build those differences into the sound. That was a lot of fun. I really enjoyed it. So that was my first sound on the DX7, completed in February of 1985.

TD: Have you been doing a lot of programming recently for your piece, working on the TX816?

JC: Yes. One of the central ideas for the piece was that I would try to get the very best piano sounds I could, because for one brief instant in the piece, the pianists will be playing sounds that are more or less within their own domain. But I took it as a challenge, because the best way to learn a system is to start by trying to simulate something. So, after the FM conga sound, I started working on piano tones. Then David Bristow came over from England, and we started working on it together. That was a very rich interaction. I think we worked out some ways using the TX816 to create some very good piano sounds. It's not a question of trying to replicate exactly, but to produce a sound where the feeling for the player and the listener both is piano-like. And I think we succeeded.

TD: Did you find yourself using the full resources of the 816?

JC: Right. For example, the various modules were assigned to different parts of the keyboard. That is a nice way to get around the 16-voice polyphonic limit in the DX series, because as you know in typical piano music low tones sustain for a long time but they're also played less often. By assigning a small number of bass notes to the first module, a few more to the second, a few more to the third, etc., you build up a system where there are beating effects, which helps the pianolike-ness of the sound. You can create a sonic form in which most of the piano literature will work, at least, and some of it quite well. The bass tones don't go away prematurely.

TD: Is it easier to get bass piano tones because there is only one string down there?

JC: I don't think so. The inherent richness of the tones is so great that we don't hear it in such a subtle way. I think that is probably the biggest reason. It's easier to make the ear think it's hearing a low piano sound. Even though there is only one string, the complexity of vibration is probably greater than in the midrange.

TD: Is that because of the wrapping and the size of the string?

JC: Yes. It's pretty mysterious. Piano's a hard instrument. The sound is so well known, probably second only to the human voice.

TD: What is your feeling about the potential use of the DX7 as a teaching tool, not only for FM theory, but for things like acoustics?

JC: I think there is a great potential. Many basic acoustic phenomena can be demonstrated quite easily using the DX7. It could become an incredibly powerful tool for learning acoustics and psycho-acoustics at a very simple level.

Beating is one whole area. Using for instance algorithm #1, turn off all the operators except the two carriers. Listen to 1, now listen to 2, now detune 1 a little bit and hear beats; and if you increase the amount of detuning, it stops being amplitude modulation and becomes kind of a rough sound.

Residual frequencies is another area. Most people who work with synthesizers think that if they hear a pitch, then there has to be energy there. There's a nice experiment you can do with algorithm #32 where you generate harmonics 1, 3, 4, 5, 6, and 7; maybe with output levels that peak above the fundamental, so that 99 is at 6 or something. You listen to them all, and you hear a pitch at whatever key you're sounding. Now if you turn off operator 1 (which is supplying the 1st harmonic - the fundamental), then there is no more energy there, but if you sound the same key you still hear that pitch. There is no energy at the pitch at which one hears it. It has to do with the harmonics and the largest common denominator, I guess. That can be very nicely explained using algorithm #32.