WIRED

Wave of the Future

By Bob Johnstone

In 1963 Max Mathews, then a researcher at the Bell Laboratories in New Jersey, published a paper in which he predicted that the computer would become the ultimate musical instrument. "There are no theoretical limits," Mathews wrote, "to the performance of the computer as a source of musical sounds." Thirty years later, you can find the father of computer music at Stanford University's Center for Computer Research in Music and Acoustics (abbreviated CCRMA, but pronounced - because this is, after all, California - "Karma").

There, in the Knoll, a 1916 Spanish Gothic edifice that was once the president's residence, with a panoramic view over Silicon Valley, a bunch of young graduate students - engineers, programmers, and roboticists, all of them also accomplished musicians - are building the ultimate musical instrument. Fashioned from software, silicon, solenoids and speakers, this virtual masterpiece will be able to replicate not just the sound, but also the feel of every piano, organ, harpsichord, and keyboard instrument that has ever existed (see "The Ultimate Keyboard," page 60). Of interest only to scholars and performers? Maybe so, but Karma's work has a way of resonating far beyond Stanford's Hoover Tower. The center's researchers have already played a key role in the ongoing metamorphosis of the personal computer from dumb terminal into multimedia machine.

In particular, they have contributed much to the development of sound boards. Over the past year or so, these plug-in PC accessories - notably Sound Blaster, the board made by Creative Technology of Singapore - have emerged from the hype as multimedia's first real-world market maker. Analysts such as In-Stat's Gerry Kaufhold predict that in 1994 sales of sound boards will top US$1 billion. Almost all (more than 95 percent) of these boards carry FM synthesizer chips made by the Japanese firm Yamaha. The chips derive from a discovery made at Stanford in 1967 by composer John Chowning, now Karma's director. They have created a revenue stream (millions of dollars in patent royalties) that has underwritten the development at the center of a new, much more natural-sounding generation of synthesizer based on mathematical models known as waveguides. This technology is now on its way to market keyboards from Yamaha, and chips from the Fremont, California-based board maker Media Vision. Because of multimedia applications, Joe Koepnick of Stanford's Office of Technology Licensing reckons that "the potential is clearly there for waveguides to eclipse FM synthesis in terms of market impact."

From New Jersey to California to Japan to Singapore and back again: for the sound of silicon, what a long, strange trip it has been. A trip that involves Phil Lesh of The Grateful Dead, if only in a walk-on part. As Karma technical director Chris Chafe cheerfully acknowledges, "It's all been totally unexpected."

When John Chowning arrived at Stanford in 1962 as a 29-year-old graduate student, he had never even seen a computer before. But as a composer, he was keen to explore the idea of speakers-as-instruments; he had encountered the concept as a student in Paris, where he attended electronic music concerts given by composers like Pierre Boulez and Karlheinz Stockhausen. So when a colleague in the Stanford orchestra passed him a copy of Max Mathews's paper describing how computers could be programmed to play instrumental music, Chowning wasted little time before heading off to Bell Labs in New Jersey to find out how it was done.

Mathews worked in the acoustic and behavioral research department of Bell Labs. There, in order to simulate telephones, researchers had figured out how to digitize speech, squirt it into the computer, then turn the bits back into sound waves afterwards. Mathews immediately realized that it would be relatively straightforward to adapt this process to the writing and playback of music. He wrote a program that made the technology accessible to non-scientists, then invited composers to come by the labs to try it out.

In retrospect, the rigmarole these computer-music pioneers had to go through in order to hear what they'd written seems agonizingly slow. As Mathews recalls, "we had decks of punch cards on which the computer scores were produced, which we would carry around in boxes." These they would load into a car, drive into Manhattan to the IBM building, on Madison Avenue and 57th Street. There, in the basement, was a mainframe computer on which time could be rented (at the astronomical rate of $600 an hour). "We would queue up," Mathews says, "then, when it was our turn, we would run down the stairs, stick our cards in the deck, and press the button." The result would be a tape full of digital sound samples, which they would take back to Bell Labs and play back through a digital-to-analog converter.

Why were composers prepared to put up with such a long drawn-out process? Because the alternative could take much longer. What was a matter of hours - compared with the several years it might take to interest some orchestra in playing their score? ("The reason I keep these expensive gentlemen with me," the late Duke Ellington once said, referring to his orchestra, "is that unlike most composers, I can immediately hear what I've written.") A second attraction of computers was that not only did they play the score exactly as written, they also offered composers the chance to go back and change bits that they didn't like. Now the challenge was how to make electronic sounds interesting, how to brighten up the dull tones to which output devices like oscillators were limited.

Chowning returned to California clutching the box of punch cards Mathews had given him. He found a place to play them at Stanford's newly established artificial intelligence laboratory, a heady intellectual environment where engineers, scientists, mathematicians, philosophers, and psychologists gathered to see what they could get computers to do. One night in 1967, while experimenting with wildly exaggerated vibratos - fluctuations in pitch often added to electronic sounds to give them a more realistic quality - fooling around with a couple of oscillators, using the output of one to control the other, half fearing that he'd break the computer if he went too far, Chowning heard something remarkable. At a frequency of around 20Hz, he noticed that instead of an instantaneous change in pitch from one pure tone to another, a recognizable tone color, one that was rich in harmonics, emerged from the machine. It was a discovery that an engineer would have been unlikely to make. What Chowning had stumbled upon, it later turned out, was frequency modulation - the same technique that radio and television broadcasters use to transmit noise-free signals. Of this, the composer was blissfully ignorant: All he wanted to do was make colorful sounds. Chowning began tweaking his algorithm and pretty soon, as he recalls, "using only two oscillators, I was making bell tones and clarinet-like tones and bassoon-like tones, and I thought, you know, this is interesting."

But who was interested? Certainly not the Stanford authorities, who, after evaluating Chowning's discovery and two of his subsequent compositions, turned down his application for tenure. Nor were US electronic organ makers - companies like Hammond. To generate its unmistakable sound (remember Booker T & the MGs?), Hammond used an electromechanical system consisting of toothed iron disks that rotated in front of electromagnets; they in turn generated voltages that formed the pitch for each key. The Chicago-based company sent its technical people out to the West Coast to check out the technology, but the engineers couldn't really see how all this digital computer stuff related to what they did. "It was just not a part of their world," comments Chowning. (Hammond went out of business in 1985; today, only the brand name remains, the property of Suzuki, a small Japanese keyboard maker.)

One of the few people who did get it, and who encouraged Chowning to continue with his work, was Grateful Dead bassist Phil Lesh. Himself a sometime composer of orchestral music, Lesh dropped by the lab for a listen one day in early 1972. Another, more significant visitor, was Kazukiyo Ishimura, a young engineer sent to Stanford later the same year by Yamaha, the largest maker of musical instruments in the world. It took Ishimura just 10 minutes to understand the principle of FM synthesis, and its potential. As Ishimura, who today is Yamaha's managing director, recalls, "We believed that this technology might be the future of music."

The reason he was so fast on the uptake was that Yamaha had already embarked on the development of digital instruments. Ishimura's boss at the time, Yasunori Mochida, envisioned digital integrated circuits - chips - as tools for making new sounds. At Yamaha's research laboratories in the small Japanese port city of Hamamatsu, half way between Tokyo and Osaka, Mochida and his team of six young engineers had tried all sorts of approaches, but without much success. "We weren't digital specialists," says Mochida, who now teaches a course in multimedia at Tokyo's Kogakuin University, "so we went looking for people who were, to ask their advice on how to make all-digital musical instruments." And, via a contact from the Stanford technology licensing office, found John Chowning, and immediately began negotiations for an exclusive license for rights to the FM patent.

"As an engineer, you are very lucky if you encounter a simple and elegant solution to a complex problem," Mochida told Music Trades magazine in 1987. "FM was such a solution and it captured my imagination. The problems of implementing it were immense, but it was such a wonderful idea that I knew in my heart that it would eventually work."

Synthesizing musical notes is a tough problem because it has to be done fast, in real time. Yamaha's current single-chip synthesizers are special-purpose digital signal processors that can zip through 20 million instructions a second, faster than most microprocessors. But back in the mid 1970s, when Mochida approached suppliers like NEC and Hitachi about making such chips, "they told us to stop thinking about something so difficult." Against strenuous opposition from the company's board of directors, Mochida proposed to Yamaha's then-president, Gen'ichi Kawakami, that Yamaha would probably have to spend hundreds of millions of dollars in order to become a chip maker in its own right. And, in true jaw-jutting corporate samurai style, Kawakami agreed, saying (according to Mochida), "if we can make the best musical instruments in the world, then no matter how difficult it is, no matter how much money it costs - we'll do it."

Turning FM synthesis from a software algorithm that ran on mainframes into chips that powered a commercial synthesizer took seven years. But from Yamaha's point of view, it was worth the effort. Launched in 1983, the DX-7, Yamaha's first mass-market implementation of FM technology, was a huge success, eventually selling more than 200,000 units, ten times more than any synthesizer before or since.

Professional musicians like Chick Corea loved the DX-7 because it had a distinctive sound, was simple to program, and could produce a variety of effects. Also, priced under $2,000, the DX-7 was affordable, and it quickly became part of every self- respecting keyboard player's set-up. Yamaha leveraged its investment in the technology across its entire product line, sticking FM chips into everything from mini-keyboards to top-of-the-line organs.

At the same time, Yamaha was looking into computer applications for FM. That was where the company blundered. Mochida decided to build a multimedia computer with built-in sound and graphics. But in a move typical of early Japanese entrants to the PC business, Yamaha tried to go it alone, developing everything, including the operating system and applications software itself. The result was a complete flop (though the project did have one important by-product: Yamaha's experience with multimedia chips won it the contract to make the sound and graphics processors used in all current Sega game consoles). Mochida was demoted, and, deciding that the chip business was less risky, Yamaha more or less withdrew from the computer market. The company did produce one soundboard - for the IBM PS/2, in 1986 - but without much support, it died a quiet death.

Today's soundboard business is largely the creation of a most unlikely pair: Martin Prevel, a French-Canadian professor of music at the University of Quebec, and Sim Wong Hoo, a young Singaporean entrepreneur. Both began by attempting to sell educational music products, but they soon discovered a much bigger market opportunity: PC game developers like Sierra Online needed sound in order to compete effectively against Nintendo. In 1988, Ad Lib (Prevel's company) brought out a board based on Yamaha's FM chip that enabled the PC to make music. But Creative Technology (Sim's company) discovered that music by itself was not enough. "It was like silent movies with a piano player," says Broderbund sound director Tom Rettig. What game developers also needed was a digital sound-output device - like the one in the Mac - to enable them to create sound effects (like creaking doors) and voices for their characters. Sim soon got the message, and the result was the Sound Blaster (see "Loud and Clear," page 62).

Voices and sound effects are created using samples, digital snapshots of sound waves that are stored in computer memory. The more sounds you want, the more space you need to store them, the more expensive it becomes. FM synthesis scored over sampling because it could generate a wide range of sounds without any memory. But though relatively rich, the sounds that FM produces are still unmistakably artificial. As memory became cheaper, and data compression techniques improved, sampling came into its own. Today, sampling - also known, confusingly, as PCM, for pulse code modulation - is the technology of choice in the synthesizer business, and many soundboard makers (including belated re-entrant Yamaha) see sample-based solutions as the logical replacement for FM. To musicians and composers, however, the technology has one serious drawback: as you would expect from sounds pasted together from frozen snippets, it lacks expressiveness. How to produce sound as efficient and expressive as FM, but offering the quality of sampling? This question drove Karma's Julius Orion Smith III to develop waveguides, the latest generation of synthesizer technology.

Smith's answering machine plays what must be one of the shorter messages around: "This is Julius...." It aptly reflects the way Smith's engineer's mind works: identifying the nature of the problem, reducing it to its essence, coming up with an efficient solution. "I'm always rating the effectiveness of everything I do," he says.

As a 9-year-old in his native Memphis, Julius Smith won a math contest. By the age of 16, he knew he wanted to be a musician. But it was not until 1980, when he arrived at Karma, that then 30-ish Smith came across the Violin Problem, a challenge that allowed him to draw on both his talents. "As a musician, I knew there were no good string synthesizers, and I thought, well it must be hard, because a lot of companies had been trying to do it for a long time." So, methodically, working 16 hours a day, Smith dedicated himself to accumulating the arcane knowledge he needed to solve the problem.

His approach was straightforward: He set out to create mathematical models of the way a string vibrates when a bow is drawn across it. Easy to say, formidably difficult to do. But in 1985, after years of banging his head against a wall, Smith finally broke through. Drawing on work done on power transmission lines in the 1920s, he recast vibration as a wave traveling in only one direction. Still, solving the resultant equations would have kept a supercomputer crunching numbers for weeks. So Smith used some fancy math to reduce by 100 times the number of calculations required to calculate the wave. Et voila: the virtual violin! It came with an unexpected bonus: since there is no difference mathematically between a violin's vibrating string and a clarinet's column of air, Smith found he could use the same equations to simulate wind instruments like oboes and flutes, too. Colleagues at Karma subsequently exploited waveguides to produce convincing simulations of other sounds. Perry Cook has developed a disembodied singing voice, a virtual diva called Shiela. Graduate student Scott VanDuyne is working on two-dimensional waveguide algorithms to create virtual percussion instruments like gongs and cymbals, traditionally among the most difficult sounds to synthesize.

In addition to versatility, another big advantage of waveguides over samples is their ability to simulate natural parameters like breath strength - how hard a reed player is blowing. By slightly varying these parameters, you can make a clarinet squeak, say, or a sax growl. And because of subtle timing issues, it sounds slightly different each time you play it - just like live music, in fact. Waveguides can also simulate howling guitar feedback, a category of sound that no other kind of synthesizer can produce.

Many of these features are included in Yamaha's VL-1 synthesizer, the first commercial waveguide instrument, which the company announced at the end of November. The $7,000 instrument drew rave reviews from the technical press: "[It's] pretty exciting," says Mark Vail, technical editor of Keyboard Magazine, "[samplers] have been around for a long time, and there's a staleness in the music industry - people have been waiting for something new to come along."

Since signing a contract with Stanford in 1989, Yamaha has reportedly had a hundred engineers working on the development of waveguide instruments, cranking out the algorithmic variations. This gives the Japanese firm a huge head start on rival instrument makers. This time, however, Yamaha does not have a lock on the technology: Its license is non-exclusive. Four US companies have already signed up to develop waveguide technology and at least as many more are interested. Leading the pack is Media Vision, which hopes to have a synthesizer chip ready for computer use in early 1994. "It's a substantial breakthrough," claims Media Vision vice president Satish Gupta, "it has the potential to completely change the rules of the game."

"Programmers are going to drool over waveguides," predicts Perry Cook, now chief scientist at the company. "They're going to want to work with this." Broderbund's Tom Rettig agrees. "To me, waveguides offer really thrilling possibilities," he enthuses. "The most exciting part is you'll be able to describe instruments that are as expressive as the most interesting acoustic instruments - and that's where current electronic technology falls down."

Max Mathews's 30-year-old prophesy about computers having the potential to generate any sound the human ear could hear may finally be coming to pass.

The Ultimate Keyboard

Under the bench in his Karma workshop, a small, high-ceilinged room that might once have been a pantry, Brent Gillespie keeps a model of the action of a grand-piano key. An intricate mechanism made of ivory, wood, felt and metal that forms a bewilderingly complex sequence of cranks, levers, springs, pivots, rollers, checks, and dampers, it provides a two-way interface between a player's fingers and the piano's strings.

The action is vital for musicians: it gives them the expressive control over an instrument required for fine performance. ("Aside from its beautiful tone, the thing that I like best about the Baldwin piano is its fantastically responsive action," reads an endorsement by George Shearing in a magazine ad.) "Synthesizers were a big turnoff for musicians at first encounter because they didn't feel right," says Gillespie, a graduate student in mechanical engineering and an expert in force-feedback systems. "My project is all about putting the feel of a grand piano back into a synthesizer keyboard."

To this end, Gillespie has built a prototype "virtual" action. A small clear plastic box from which two keys stick out, its sensor keeps track of the position of a key as it is pressed; a solenoid puts out an opposing force proportional to the key's displacement. It's uncanny: you press the key and you feel the striking of a string that you know is not there. The box can be programmed to replicate the different feel of instruments as similar as pianos and harpsichords, whose strings are plucked rather than struck.

Why hook up a waveguide synthesizer that can reproduce all possible keyboard instrument sounds? John Chowning explains: "We have a generalized keyboard that can be particularized to any desired piano or any specific piano. If you want a Yamaha, you can have it. If you want a particular kind of feel on your Yamaha, you can program the resistances. Or if you want a forte piano of, say, the 1780s, you can have it, and the sound that goes with it.

"We have this idea of a piano which in all essential respects - auditory, kinesthetic, tactile - is a piano, only it has no strings, no action. But it supports the repertoire for which these instruments exist. It's easy to keep in tune, and you can easily change the tuning system from, say, mean tune, which you might want for the 18th century, to well-tempered, as in Bach, to equal tempering, as used today, just with presses of buttons. You can play it at night because you can turn off the loudspeaker and listen through the head-phones - that's important in Japan. And it's easy to move. It's the ultimate piano.

"And we have an historian, George Barth, whose scholarly expertise is in the evolution of keyboard instruments. If you have a replica of, say, a 1780 forte piano built, with an artisan who does it cheaply, it costs $20,000, and with artisan paid at normal rates it would cost $100,000. George Barth has one, but what do his students do? Well, they have to convince their employer, or their university, or if they want to perform, they have to come up with $100,000.

"This is the general solution, you see, for the extension of scholarly activity - and it really democratizes the idea of performance. No longer will it be true that only the wealthy kid gets the good Steinway, but every kid gets the good Steinway."

Loud and Clear

If ever a person was in the right place at the right time, it was Sim Wong Hoo, chairman and CEO of Creative Technology, and one of multimedia's first multimillionaire. The place was San Francisco, the time August 1988. The 32-year-old Sim had arrived in the US from his native Singapore to market his company's pride and joy: the Creative Music System, a synthesizer card whose software enabled users to compose music on the PC. A group of potential customers for the system was Bay Area-based game developers. But as Sim went 'round talking to the companies, he quickly realized what people really wanted not just another music synthesizer, but a board that could handle digitized sound, to enable the PC to produce sound effects and speech. "Sim had a clear vision of the importance of audio, at a time when the industry had just started with it," recalls Tom Rettig, sound director at top educational game maker Broderbund, "he contacted us at just the right time."

Sim's vision had deep roots. "I felt that computers should be more human-like," Sim says, "able to react, to talk, sing and play music." In the mid-1980s, Creative designed a series of computers for the Singapore market featuring rudimentary (Chinese) speech capabilities. But as competition in the clone market became fierce, Sim switched focus from PCs to the add-on board business, where profit margins were higher.

In 1988, the tiny market for sound boards was dominated by Ad Lib, a Quebec-based company whose Yamaha FM synthesizer-chip-based board was supported by hundreds of game titles. At the time, Ad Lib was the only firm that Yamaha supplied. Then Microsoft stepped in and asked Yamaha to sell the chips on the open market. The Japanese firm agreed. Creative's great good fortune was to be the first to come out with a board that mounted the Yamaha chip - making it compatible with existing games - and that supported the new software. Sound Blaster was launched in November 1989. In addition to music synthesis, the Sound Blaster also offered the digital sound capabilities of the Mac. "That combination really made the whole thing take off," says Rettig. Broderbund developed two of the first products that supported Sound Blaster: Princes of Persia and Where in the World is Carmen Sandiego? Developers appreciated Sim's aggressiveness and his willingness to put his company's technical resources at their disposal. If they needed a software driver, for example, Sim could conjure one up for them overnight, making good use of the 16-hour time difference between Singapore and the West Coast. (Business hours on the island begin just as the US work day ends.) An Asian manufacturing base enabled Creative to drive down board prices, giving it a competitive edge that Ad Lib could not match. From an initial listing of $299, the price of Sound Blaster eventually dropped to below $70, as the market for the board exploded. In Sound Blaster's first year, Creative sold 100,000 boards, a phenomenal amount for the time. Today, the company is on a roll, selling 300,000 boards a month.

Now Creative's goal is to branch out from its beachhead in audio to stake out other parts of multimedia, like CD-ROM upgrade kits and video boards. "I'm not stopping here," Sim says.

Bob Johnstone is WIRED's contributing editor in Japan.