

C00066 00014 BILL SCHOTTSTAEDT: COMPUTER COMPOSER'S COMPOSER  
C00073 00015 A TALK WITH CHRIS CHAFE  
C00086 00016 BIBLIOGRAPHY  
C00089 00017 DISCOGRAPHY  
C00091 00018 THE AUTHOR  
C00093 ENDMK  
C^V;  
WATCH YOUR KARMA

by Stanley Jungleib

It's the Stanford Center for Computer Research in Music and Acoustics (CCRMA). It's one of the two most influential computer music centers in the world. It's called "Karma" for good reason. And its last name is "Chowning". It's his baby.

((photo Chowning))

John Chowning played violin as a child but switched to percussion as a teen. During the Korean war he studied jazz at the Navy School of Music, then went to college to study composition. From 1959 to 1962, he studied in Paris with Nadia Boulanger. Here he was also exposed to the European electronic scene, hearing works of Boulez, Stockhausen, Berio, and others.

In beginning graduate study at Stanford in 1964 he found no electronic music and little interest except on the part of his composition teacher, Leland Smith. With help from the Stanford Artificial Intelligence Laboratory, Chowning installed Music IV software provided by Max Mathews of Bell Labs, and this system made its first sound in September, 1964.

In 1966 Chowning joined the Stanford Music faculty. With the system up and running as part of the AI lab, located in a strange and lonely building far removed from the main campus, Smith and Chowning began teaching computer music. Smith worked on the music composition interface and input language. His SCORE system for computer music became a standard in most computer music centers. Chowning worked on synthesis techniques, timbre, reverb, and the Doppler effect (the pitch shift which occurs when vibrating objects approach, then leave you).

Chowning's 1972 composition "Turenas" was a virtuoso demonstration in sound programming using frequency modulation (FM) techniques which he had "stumbled upon" in the mid-60s. It immediately became a classic and put Stanford on the computer music world map. Chowning's paper describing the FM technique appeared in 1973, and Stanford patented it. Ten years later, Yamaha introduced the highly popular DX-7, which implemented FM in low-cost, LSI components.

As the computer music project's reputation grew, students

gathered. And with increasing demands on the system, it became clear that the project required its own computer and synthesis resources. Design work began on what became the "Sambox" research synthesizer. They also consulted with Boulez on the design of IRCAM in Paris. CCRMA was formed in 1975 with John Chowning, John Grey, Andy Moorner, and Loren Rush as co-directors. CCRMA and IRCAM have served as models for all other computer music facilities in the world.

In 1977 the Sambox was installed. In 1979 the AI lab left the building to Computer Music. It was at this point that CCRMA in fact became a Center, with Chowning as Director. The environment forged a close-knit tribal/familial atmosphere. Chris Chafe raised two of his children there. They began a tradition of quarterly outdoor concerts which stimulated interest in the center, inspired and attracted students, and developed a certain mystique within the generally conservative Stanford community. In 1980, CCRMA received a three million dollar grant from the System Development Foundation, which helped establish the current activities of the Center.

#### CCRMA TODAY

Today, Chowning and Administrative Director Patte Wood keep CCRMA vital, free, relevant, and unpretentious. Perhaps the greatest tribute to their work is the sense that you get from talking with the staff and students, that even if they were working with slide rules and kazoos, everyone would still want to be here. The Center takes great pride in its cooperative, international, and interdisciplinary environment which supports ideas that might not arise in other, more structured settings. Chowning has also just added a great deal of historic authority to the center in the personages of his former mentors Max Mathews and John Pierce, from Bell Labs days. Current support is from the National Endowment for the Arts (NEA), National Science Foundation (NSF), Doreen B. Townsend, the Rockefeller Foundation, and private gifts.

CCRMA activities can be roughly classified as composition or research, but there is plenty of overlap between the two areas, since most of the composers do research, and most of the researchers compose! A great deal of the energy of the center supports composition. It is an art facility, committed to new and future music worthy of the instruments they are privileged to be investigating. Works created at CCRMA dominate the computer music genre. CCRMA celebrates virtuosity in this area through its regular concerts. But none is like the outdoor summer concert at Frost Amphitheatre, a gently hollowed bowl of ground. The Grateful Dead play here every spring. And CCRMA holds the cerebral counterpart: You lie still on the grass, watch the stars, hear bizarre FM patches jetting across quadraphonic space, and think, "what a long, strange trip this electronic music business has been!"

There are about 40 people active at the center, including professors, the research staff, composers, graduates and

undergraduates. CCRMA research covers many areas: digital synthesis, signal processing, interactive composition, MIDI performance, graphics, machine recognition, and psychoacoustics. The research going on today and described later in this article will have an immense impact on the equipment you will be working with in the next decade.

## INDUSTRIAL RELATIONS

The Center's reliance on powerful computers is an expensive habit. So, the frustrations of grant-seeking come with the research territory. Except where there may be a discovery affecting robot perception, the federal government is not that keen on computer music. After all, you can't bash Russians or Iranians with it. They seem to be willing to let the commercial sector sort it all out. As a result, there have always been ties between CCRMA and the music industry.

For example, Chowning has worked with Yamaha since 1967. When you buy a piece of Yamaha FM gear, some of your cash ends up in Stanford's pocket. They say it is not that much because the royalty is figured only on the FM chips, and that these are a small percentage of the overall cost of the whole machine. But consider the quantity... Yamaha is very well represented at CCRMA with their power amps in almost every room, lots of DXs and modular FM tone generators, several sequencers and drum machines, plus an entire MIDI studio (KX-88, TX-816s, SPX-90s and so on). The only non-Yamaha MIDI hardware around are the Sequential Studio 440 and Opcode Mac interfaces.

As another example, Passport Designs President Dave Kusek says "I searched literally the entire world for the best music printing technology I could find." And he found it at Stanford in Leland Smith's SCORE. Smith has been perfecting his professional, engraver-quality manuscript system since the early '70s. Passport has licensed Smith's program, added a mouse-assisted user interface, MIDI input in step mode, and four tracks of real-time MIDI playback. The package is to be released for the PC this September, at \$795. Kusek describes demand as "absolutely frightening."

To profit from even more commercial products, CCRMA is actively inviting research support and has begun a new program to organize its industrial associates. Unfortunately, some of the more speculative research that they are engaged in may not pay off within the time span that a small company can afford. So, we may see only big players in this game.

May 27-29, 1987 was the First Annual Meeting of the CCRMA Industrial Associates. Participating companies were treated to a three-day tete-a-tete with the staff, students, and each other. By getting everyone together at once for an annual presentation of all the work, everyone learns more and the discussions are more productive. The first day presented major conclusive results of recent work. The second focused on more speculative areas in process. The third identified promising trends CCRMA

thought especially deserving of a commercial response. The presentations were punctuated with fine conversation and seasoned observations from notables such as John Pierce and Pete Samson. By all accounts, it was a roaring success, leaving everyone satisfied that there had been genuine communication and exchange. Seeing how this program boosts commercial development of CCRMA research should prove quite interesting.

#### THE KNOLL

CCRMA moved into its current facility at the Knoll in June, 1986. Built in 1916 as the home of the University President, it looks every bit the part. The Renaissance-style mansion is built into a hillside so that from the front you see an impressive and imposing three stories, but in the back you get a cozy courtyard bordered by only two stories. It may be Stanford's classiest structure. Its location on the Knoll gives it parental watch over the rest of the campus. The southern view looks down on Silicon Valley, filled with Stanford graduates. Some prefer the western view: the seasonal Lagunita lake, oak-studded foothills and Pacific sunsets. In any case, it's a damn fine setting.

((photo Knoll front))

((photo Knoll back))

Inside, the ornate woodworking has been de-emphasized by sedate, creamy walls, with warm white ceilings that reflect the incandescent lighting. There are most of the original lamps and mirrors. And just a few well-chosen computer music posters. Otherwise, no distracting wall graphics. Unlike modern office buildings, you can actually open the windows. The kitchen includes an espresso machine--which is always hot.

The building has two separate recording facilities. Upstairs is the performance-size ballroom, with its control room behind it. This control room contains all dubbing equipment: audio output from the Sambox, tape machines in various sizes, Beta players and PCM converters, CD players, and so on. They have to be able to handle all formats. There are over 400 audio channels in the room. An Apple II programs all of the routings between the audio peripherals. The second recording area is downstairs. There is a very quiet, isolated, medium-size recording studio with adjustable acoustics. Its accompanying control room has been used for digital mastering and will eventually be a digital multi-track facility for producing compact disks.

#### THE SYSTEM

The Knoll has its own computer network, which has been maintained for many years by the wizard Tovar. CCRMA is unimpressed with UNIX, preferring their WAITS operating system, which has its origins in the Stanford AI lab. CCRMA's composition and digital signal processing

(DSP) libraries are probably the most complete in the world. It therefore doesn't pay to switch operating systems around just to be fashionable. When more or different computing resources are needed, Ethernet connects the Knoll to the rest of the Stanford network (LAN) and, for that matter, to the rest of the world.

The central processor is the cult-favorite Foonly F4. (It is interesting to note that Foonly, Inc. includes David Poole, who was with the AI project and assisted the young Chowning.) The mainframe controls 16 graphics terminals (which are spread around the building), the disk drives, the Sambox, and the audio switcher which routes Sambox output to the various work rooms. The drives use 300-Mbyte removable packs so that memory space can be shared. The disks are used for programs or sound files: at a sampling rate of 44.1 kHz, you can put about 23 minutes of stereo sound on one.

There are three main types of music systems at CCRMA. First, most teaching and composition, and some research, is done on the legendary Sambox synthesizer. Second, advanced research in acoustic analysis and pattern recognition using artificial intelligence, and physical modelling, is performed on LISP machines. There are a half-dozen of these. Third, MIDI research is supported in its own studio and at workstations in various places.

We will now look more closely at how each of these three types of systems are being used.

#### THE SAMBOX

CCRMA's central teaching and research tool is the Systems Concepts Digital Synthesizer designed by Peter Samson, nicknamed the "Sambox." It was built from 1973 through 1975 to be "the biggest and most powerful digital music synthesizer of its time." Its goal was to have "a real-time synthesis capability comparable to that of a string quartet, both by additive and subtractive methods." Since its installation in 1977 it has attracted a steady stream of awestruck students and composers to CCRMA, and has played around the clock to the toughest audience you can imagine. All give it the highest praise. Chowning calls it "the Stradivarius" of digital synthesizers.

The Sambox is a fully-programmable, general-purpose synthesizer. It is a very intelligent peripheral to the main system. Depending on the sampling rate you select, up to 256 oscillator/generators are available. There are 128 modifiers (arithmetic operators) available for modulation, mixing, resonance, and so on. And 32 delay units, which of course form the basis for reverbation, but can also be used for digital filtering. "General-purpose" means that the composer can configure the Sambox's resources for different synthesis techniques, as needed in a composition. This is done using a set of algorithms, called "instruments." You can simultaneously use additive synthesis (with different envelopes for each harmonic), FM (which uses much fewer oscillators), and subtractive filtering (using the modifiers for resonant filters). In addition, each instrument usually has a wide range of

delay or reverb effects. There is no physical, real-time input. The sound that comes out is determined solely by how you choose and calculate instrument variables in your composition program.

Originally, audio output from the Sambox always went to all the terminals. You always heard everything that everyone was working on. Likewise, your own work was always on stage. This constant interchange of ideas and direct musical communication is part of the reason for any "CCRMA sound" that may exist. But it also inspires a subtle competition. There is something humbling about there being a lot of good music piped around the building all the time. It's like trying to fly with the Blue Angels of computer music. Who wants to be heard making trivial sounds with such a great resource? Now, by default, the mainframe's audio switcher routes Sambox output only to the composer's terminal. But you can easily defeat this routing and still hear everything that is going on.

Also, the Sambox originally had four audio output channels, and a lot of CCRMA composition has been done for quad space. Now they have added four more channels. Some of the drier rooms will soon get eight speakers so that you can do three-dimensional room simulations, or draw 3-D sound paths. There are a whole new set of non-trivial challenges in synthesizing the movements of the player in vertical space, the directionality of the bell of a horn, and so on.

#### PLA-ING IN THE SAMBOX

To start composing at CCRMA you first learn to write a basic score with the PLA composition language, pass it through the PLA program, which converts your score into a "note list," then pass it through a compiler which converts the note list into Sambox machine instructions. Sometime later, you feed these instructions to the Sambox, and it computes all the sound called for by your program, in real time. Because the Sambox plays everything, the main computer system is not interrupted or slowed down.

((illus: a simple PLA program))

But notice that you have no real-time interaction with the Sambox. Once you set it on its way, the Sambox is off and running. Coming from the MIDI performance world, it is hard to switch over to this detached work rhythm: changing a score, waiting for it to compile, then hearing the result later (sometimes, hours later). I was reminded why performance equipment was invented in the first place. For example, my seven-minute piece soon began taking an hour and a half to compile. While waiting, you can strategize about working in smaller units, make coffee, or read a spy novel (the method of speech synthesis researcher Bob Shannon). But if you start another computer job, you just slow down your main job.

The other thing I found frustrating is that until one learns

how to allocate the Sambox operators efficiently, novices like me run out of generators quickly--there may not be enough generators available to play a note that is requested by the score. Genuine composers at CCRMA (including Schottstaedt, McNabb, Chafe, Wolman, Berger, Malouf, Karpen, Taube, Krupowicz, and others) know how to cheat this phenomenon, and are able to achieve wonderful densities.

Why go through this hassle? Actually, it's a small price to pay for learning a completely different way of thinking about music. And, of course, as you learn more about it, you begin to understand and appreciate the skill and hard work it takes to realize a convincing piece. The power of the Sambox is basically that any note can have its own programmed set of parameters. You therefore have complete and arbitrary control of timbre over time. You can set your voice parameters to fixed amounts, for specific sounds, or you can calculate them depending on any number of variables which you might choose or define. Your timbre and time controls are completely integrated within a single set of consistent tools and expressions.

((start Schottstaedt sidebar))

#### SAMBOX TIMBRE RESEARCH

In addition to supporting composition, the Sambox is still the seat of several promising analysis/synthesis projects which may result in new types of synthesizers.

One that has caused a big stir is the Waveguide Filter (WGF) technique created by Julius Smith, an authority on digital audio filtering. A waveguide is a software delay line (similar to your DDL), which also happens to be bidirectional, as well as lossless. To create filters of varying complexity, you interconnect waveguides of various delay times. This gives you rich, active filtering. By inserting some strategic losses you can easily make a great, complex reverberator. Or, if you close the network, you can make an oscillator with a very complex and dynamic timbre. The example I heard sounded like Darth Vader being short-looped at the bottom of a well, which only means that I really can't describe its strange impression of constant timbre in constant motion.

Because musical instruments, especially strings, woodwind bores, and so on, are nearly lossless, waveguides are much more efficient than earlier synthesis techniques. By introducing signal-dependent valves and junctions in the waveguide network, you can generate interesting non-linearities (transients) in essentially the same way as they are generated in real instruments, at a very low relative cost.

((illus waveguide network))

A variety of people are still trying to stretch the limits of

additive, FM, and subtractive techniques. Composer Mike McNabb works on timbre modulation by dynamically mixing additive-type waves. A variety of synthesis methods are being combined to produce realistic bar percussion or marimba tones (by Xavier Serra), piano tones (by Joe Marks and John Polito), or a variety of instruments (by David Lo).

When a machine in the Sambox's position reaches ten years of age, questions naturally arise regarding its usefulness. Maybe it's time for the gold watch and testimonial banquet?

No chance. The Sambox is still vital. It is getting as much use as it ever has, and supports some interesting work. The general impression is that there is a lot of music left in it still untapped. There are a number of algorithms still available that haven't yet been used in any pieces and are probably interesting.

At the same time, even its creator writes, "In an increasing number of cases, its capabilities are close to the minimum needed for serious work." So, it is possible to respect it, while still looking forward to the next generation of decentralized workstations. The Sambox's main drawback is that you can't really control it in real time. You can't feed it performance data and have it alter its sound in response. (It took performance synthesizers to provide velocity control of FM.) The cutting edge of research requires much faster, real-time systems, for exploring programs that display musical intelligence. This brings us to LISP and the drive towards musical expert systems.

((recommended article split point))

#### ARTIFICIAL INTELLIGENCE FOR MUSICAL ANALYSIS

CCRMA makes good use of its several Symbolics and Xerox LISP machines. They like them because programs are easy to modify in response to developing ideas. The LISP machines are used in three main ways: to study music analysis and recognition, as synthesizers for physical modelling of bowed strings, and as high-powered MIDI controllers. A main advantage of the LISP environment is that it allows all of these different applications to be easily interfaced to one another.

There is one problem at the root of everything, which has received a lot of the Center's attention since 1981. How can computers be made to perceive music? If you sample music at any point, you'll see wave energy at various frequencies throughout the audible range. In a medium as unpredictable as music, how do you connect these timbre pictures together and infer the existence of specific musical events? How do you know when a frequency is a fundamental or a harmonic? How do you untangle a mass of sound into the structural forms behind it?

We don't know yet. These questions are enormous. Though some promising results have been achieved, current work



is quite experimental. They are still defining the problems.

And what are the rewards? A polyphonic analyzer would obviously be useful for music transcription: creating a musical score from a performance. That means identifying the notes correctly, correcting any mistakes of the interpreter, and establishing the correct musical context: key signature, time signature, with all the correct notation.

Another application would be for computer assistance in editing digital audio. An intelligent editor would allow us to point to a score and give a command, and it would find the sample and implement the edited version of the passage.

Third, analysis is essential for tracking live performance. Acoustic analysis is a way around costly and problematic physical sensors.

Fourth, improved analysis of performed music allows us to better understand nuances. If we can place performance under a microscope, then we can learn how to emulate the control that a real performer exerts, within a synthetic system.

The analysis problem tends to divide into: 1) converting acoustic input into physical events, and 2) converting physical events into musical elements. The first area has been investigated by Chris Chafe and David Jaffe. Bernard Mont-Reynaud works on the second area.

((start Chafe sidebar))

How does artificial intelligence fit in to the question of musical perception? Let's assume that you are trying to track pitch. In pitch detection you have some pretty straightforward bandwidth/time tradeoffs. With linear signal processing tools (FFTs, autocorrelation), you can't tell what the period of a 100-Hz tone is until you have a couple of periods, which takes 20 milliseconds. So you are always a little behind.

The question is, why doesn't the ear have this problem? It doesn't need to wait to tell what is going on. In fact, we can analyze long waves and rhythms quite quickly. How does it do this? The studies coming out of neurophysiology are pointing out that hearing is capable of being massaged according to information provided from contexts that have been constructed at a higher level of perception. That is, when you are listening to a musical line, it is quite likely that your ear is actually tuned to be a very responsive filter to tones in the region that your mind expects. (This is roughly the same mechanism that allows the ear to quickly identify different musical styles, or focus in on one voice in a crowd.)

If hearing is active, rather than passive, then you need to put tunable receivers or windows in your musically sensitive robot. This is where the artificial intelligence

comes in. If the processor has some clue about what the musician is about to do, if it is already near the right place, then it will have an easier time keeping up. Of course, music is full of broken predictions, so this is, at best, a nasty area to get into. But the research has shown that you stand a much better chance using a system that knows how to change its analysis strategy according to the specific situation. The system keeps in its artificial mind a model of the timbre it is supposed to be analyzing. With this information, it looks at the input in three different ways, to detect different classes of events. As the system builds a processing history, it uses successes to further tune its perception.

If acoustic analysis is the input function, then the output function of the polyphonic processor is a graphic display of the score for the input. Mont-Reynaud has demonstrated some remarkable examples of acoustic music transcription. The initial work was actually done on the Foonly (in SAIL) and achieved the power to transcribe both metrically and melodically a monophonic line of moderate complexity from Bach or Mozart, or a drum part, with no restrictions on variation of tempo. The basic technique proceeded by performing a low-frequency FFT analysis to detect the occurrence of most fundamentals, then applying pattern recognition techniques in successively smaller time slices. The analysis addressed musical problems: meter, tempo variation, note values, key and modulation, accents, and so on. The system eventually achieved the rather amazing ability to transcribe music of the complexity of Joplin's "The Entertainer," for piano.

((illus: Mont-Reynaud's chart of music analysis techniques))

For the next stage of work, Mont-Reynaud has shifted to the LISP machines. The system under development displays a high-resolution spectrogram of complex musical events. By using a "harmonic cursor," the user can suggest ideas to the analyzer about how to decode the signal. Allowing some human intervention considerably extends the range of the analysis system and improves the results.

((illus: spectrogram))

In the meantime, work continues on developing pattern recognition that will ultimately be able to decode music that is highly syncopated, non-metric, or atonal. Here again, the intelligent analyzer adapts its behavior to the input. It tries to decode a harmonic series as the overtones of a single instrument. It then considers other notes, and asks each note to "vote" on a pattern of interpretation. Under this scheme, an event may end up with high "peer approval" and will be considered part of a specific voice, or the event may be judged "anti-social" and be ostracized from the interpretation.

#### PHYSICAL MODELLING

The second main type of work done on the LISP machines

at CCRMA is in the new synthesis technique of "physical modelling." This is done mostly by David Jaffe, Julius Smith, and Chris Chafe. The modelling research creates very compelling simulations of orchestral instruments. It depends on a combination of waveshaping synthesis (using one wave to distort another), a specific processing algorithm, and lots of digital filtering.

Usually, synthesis gives you a Fourier/Helmholtz-type model to work with, such as the harmonic series, and you are asked to build your sound indirectly, using oscillators, filters, and so on. In physical modelling, complex waveforms are created by a powerful and compact mechanical simulation created from a network of linear and non-linear elements. The linear ones serve as the resonant instrument string and body. The non-linear elements control the interaction of the impulse generators (the pluck or attack stuff) with the resonating element. You therefore get direct control over the instrument by using physical parameters. It turns out that the model is especially good for plucked and bowed strings, but can be extended to simulate all reed instruments, brasses, and possibly flutes and pipe organ. Imagine a synthesizer interface where you select your instrument class, then build the initial instrument by specifying the physical and mechanical properties of your favorite acoustic instruments. When you apply a musical score to the model, the sound contains all the life that it would have if someone were playing the physical instrument. For example, the effect of bow velocity, embouchure, or hammer hardness are accounted for and directly synthesized into the sound.

There are three theoretical cornerstones to CCRMA work in physical modelling. The first piece is a computational model for vibrating strings suggested by the English mathematicians McIntyre and Woodhouse in 1979. They were concerned almost exclusively with the bow-string interaction. They didn't listen to the model or work on sound quality.

Second, is the "plucked-string" algorithm published in 1983 by Kevin Karplus and Alex Strong. As graduate students in the Stanford Computer Science department, they found a cheap and flexible synthesis technique by starting with a randomized wavetable, and computing a decay by averaging successive samples. The idea of a simple averaging was the most reasonable computation you could ask of an eight-bit microprocessor in real time.

The third element is the extensions to the "Karplus-Strong" algorithm also published in 1983, by David Jaffe and Julius Smith. Jaffe knew Alex Strong, and was interested in obtaining high-quality guitar-type sounds for composition. Smith recognized that the Karplus-Strong algorithm was a special case of the McIntyre and Woodhouse model. McIntyre and Woodhouse had apparently hit upon the same fundamental structure which describes traveling waves in a digital loop. Jaffe and Smith applied a variety of new signal processing techniques to produce high quality sound and to cover a greater range of instrumentation and expression.

Modelled sounds have a strong character and identity. You

have flexibility over the construction of the instrument: the basic type, string length and stiffness, body material, and so on. For example, using the model, you can eventually build a bigger and more complex acoustic instrument than you could ever physically build. (A 25-foot long grand piano, or Harry Partch concoction should be no problem.) In playing the instrument, you also have great performance flexibility. Instead of implementing inflections indirectly (by adjusting a filter parameter, for example), you can program playing styles like dynamics, articulations, glissandi, slurs, trills, picking direction, and pick position. All of these things are in the model. To hear an example of the string in action, listen to Jaffe's "Silicon Valley Breakdown."

As a cellist, Chafe could not ignore the bowed-string sounds he heard Julius Smith producing with the model in 1982, and immediately became interested in physical modelling. Chafe now works on controlling the model. He is looking primarily at the control functions that a string player uses during a note. These are much more complex than piano control, so progress in this area will no doubt contribute to making keyboards and controllers (wheels, velocity, pressure, pedals, breath, and so on) more expressive.

Chafe's bowed-instrument model uses a separate synthesizer for each string, and cross-couples them, as they are physically linked. The software reads an input score and makes changes in bow velocity, bow pressure, bow position, string damping, and so on, as would a real player. In the bowed model, you have to worry about what happens when the person lifts the bow off of the string. Chafe's model switches from a driven delay line to a ringing delay line. As a result, all of the effects that are likely to result from the player's actual physical movements are programmed into the sound. The software is essentially an expert cellist, that creates time-varying envelopes for the instrument controls. Variety and reality in phrasing is implemented by making lots of small changes to the instrument parameters.

All of the research comes together when you use one LISP station to fully analyze polyphonic input, interpret it as a physical model, re-synthesize the input with complete accuracy, and route realistic control output to MIDI destinations. For example, you could map violin bow velocity to your amplitude envelope and map bow pressure to your FM index. Just that would give you a good, intuitive-feeling controller for different bow articulations. That would be the first step to a really interesting MIDI-controlled string instrument. And any of these techniques are applicable to other instruments.

## MIDI

Almost as remarkable as the scope of CCRMA's basic research is the degree to which they have embraced MIDI. Originally, the computer music community was not that enthusiastic about MIDI. And "MIDI-bashing" continues in some quarters. Some think MIDI is already technically

obsolete, and many academic composers feel that performance synthesizers are not flexible enough for them to want to bother composing with. Those are reasonable concerns. But some academics are just mad because they weren't given the chance to take years researching MIDI to death through a beauracracy of study groups. MIDI was a limited, grass roots campaign fought in the streets of commercial reality. Even with its compromises, there were several points at which the original specification effort could have fallen apart. Its victory is based on its open architecture, universality, low cost, great sound potential, and real-time control. John Pierce, with the perspective of one who should know, has called MIDI one of the greatest inventions in the history of music.

His fellow Carmanians evidently agree. They have the Yamaha-stocked MIDI studio mentioned above, and they are assembling a variety of MIDI workstations running LE LISP on Macintoshes. They have stated they intend to "influence the music industry's understanding as to where available power is best applied, and participate in the definition of future technology." And, it turns out, a good MIDI facility also helps attract good students.

On his LISP machine, Mont-Reynaud has configured a MIDI system on which you can play any figure or section on a MIDI keyboard, and the notes instantly appear in score form and can be edited.

Chafe also sees new composition techniques stimulated by the appearance of CCRMA's first performance sampler, the Studio 440. The piece he is now working on uses a LISP-running Macintosh to control the Studio 440 and some FM boxes, to take advantage of the contrasting sounds. By using only sampling, you get accurate sounds but you don't get to take them anywhere. In FM you have a harder time getting accurate, detailed sounds, but you have more control over pushing them out into other timbral dimensions. By quickly splicing sounds created with the two methods, interesting effects turn up. He did this on the Sambox before: "I generated the score by creating a musical Osterizer. You stick in sampled sounds and FM sounds into the Osterizer and swirl them up. Playing with those swirl patterns was the root of that piece." And now he is working on getting these same kinds of splices over MIDI, using one of each kind of synthesizer.

Keep your eye on CCRMA. This is not your average garage band.

--Sidebars--

BILL SCHOTTSTAEDT: COMPUTER COMPOSER'S COMPOSER  
((sidebar 1))

((photo Bill))

Computer composition requires outstanding intuitive and logical skills. Having to fully analyze the inner processes of a complex and imaginative musical statement and render it to an algorithm understood by a

computer language is...well, it's real hard. PLA is the dominant composition language at CCRMA. It makes it much easier to understand and write the algorithms necessary to compose computer music. And it was invented almost eight years ago by Bill Schottstaedt, who has also done pioneering work in FM. There is a file you can study that is organized as bells, ordnance (explosions), glass, metal, and wood percussion, strings, 50 species of North American birds, insects, frogs, and voices -- and almost all are done with Bill's famous "FM violin" instrument. Since 95% of CCRMA's compositions rely on FM, many composers start with Bill's instruments and patches. Through PLA, his FM work, his unrivalled command of the Sambox, and the strength of his compositions, Schottstaedt exerts a powerful influence on all composers at CCRMA.

As a composer, Schottstaedt is oriented towards finely crafted sound sculpture, as opposed to improvisation. He is attracted to computer composition partly because it gives him pitch and timbre control that he can't get in any other medium. Schottstaedt is not interested in the performer. He asserts that pitch control is the privilege of the composer, not the performer: most players have inaccurate and idiosyncratic pitch control which renders the subtleties of alternative tuning systems unimportant. Schottstaedt also complains that a performer's movement distracts from the music. He prefers fixed sound sources with minimal, if any, reverb. His latest work rejects the quad sound system in preference for a single monophonic channel. Even titles distract from the music. What is going on here? I see an essentially Wittgensteinian aesthetic; what is visible is logical, unequivocal, and completely purged of cheap sentimentality and vague intentions. Its specificity in a sense validates the rich expression which appears on the flip side of its musical coin. In other words, formalism doesn't prevent Schottstaedt's music from striking many people as being perhaps the most beautiful and effective computer music being done today.

After not hearing much computer music for several years, I was privileged to hear the world premier of Schottstaedt's "Leviathan." It knocked me on my ear. Here was a vibrant, rich tonal tapestry several generations beyond the "blips and bleeps" by which I had dismissed the genre after an earlier exposure. But, more important than my reaction are the music concrete techniques that Schottstaedt used in this piece. They are a great lesson in the art of sampling.

All of the sound material for "Leviathan" was taken from a standard sound effects LP. Fourteen noisy sounds were collected: train, rooster, ship's mast, milk machine, water pump, thunderstorm, and so on. In the right hands, this much noise means that you can pick out any pitch that you like. Schottstaedt wrote a program that linked over 15,000 splices of these samples, and controlled the looping for each segment. To keep the sample loops interesting, he randomized the loop start and end points so that each loop would play a slightly different sample. He also uses dynamic loops. For example, in one particularly bone-chilling segment, he takes the sound of a creaking ship's

mast, and slides a short loop down it to create a modulated timbre that been described as "a compound fracture grating on ice."

Schottstaedt sometimes complains that he has taken the Sambox as far as it can go. But he keeps coming up with new things. A post-PLA system called JETSAM integrated more functions than PLA. Now he rejects both, saying that you don't need a composition language, you only need a good computer language. And which language is the prime candidate? LISP, of course. Bill just used it to create a rule-based expert system which writes species counterpoint in any standard mode for up to six voices.

A TALK WITH CHRIS CHAFE

((sidebar 2))

((photo Chafe))

As Technical Coordinator of CCRMA, Chris Chafe participates in much of its research. His background includes playing in the big band with Cecil Taylor at Antioch College, and composition study at University of California, San Diego.

SJ: There is a great deal of interest in MIDI here.

CC: MIDI is providing things that people have been looking for for a long time. I think of the Samson box as a very, very fine studio machine that permits you to assemble music in a sculptural fashion. It reminds me of a sculptor's studio where all the work takes place in one chamber and the pieces that come out of it are very finely crafted. The Sambox system is in fact unique in being so tactile. You are always perfecting the finished result. Composers aren't so used to that. Traditionally you have been separated from the sonic integration of media and score.

MIDI is an enhancement of this approach, as well as a departure. Now we can capture important performance skills and use them electronically. And we can also get out of the sculptural mode and jam. We can improvise, which we couldn't do with the studio machines.

We've broadened the pallet. Both worlds have their characteristic traits and advantages, with little overlap. Perhaps some of the workstations that are looming on the horizon will allow people to capture the advantages of using the two together.

There's a lot of good stuff you can see coming down in the next two or three years. Some of those DSP chips that are going to be available have an amazing potential for sound generation.

SJ: What basically is a DSP chip? A lot of multipliers and dividers...

CC: ...and hopefully, high-speed memory access and genuine general-programmability, as well. We can write the programs to do the work now, but the interesting stuff doesn't yet have an implementation to allow it to run in real time. It won't be an instrument in either the MIDI sense or the Sambox sense until it does run in real time. That's the thing.

SJ: You can basically do a Sambox with a couple of these DSP chips?

CC: (nods) Yeah. Right. A handful. Of course there is no free lunch--if you keep calling for greater textures, you'll eventually leave real-time, or sacrifice fidelity.

But if these things appear in systems that are open enough, we are going to have people writing programs that can run in real time and do algorithms that we can currently only hear in the form of test tones that have crunched overnight. It would be really fun.

And the same goes for analysis. We want to do real-time tracking to provide polyphonic pitch detection on stage, to have a bunch of live players and let the machine play along. Currently, we are just scratching the surface of the polyphonic analysis problem, and if our experiments are any indication, it is very compute-intensive to do that stuff. So, VLSI is going to have a huge impact on that. We are going to need to get these things to run very, very fast.

SJ: I had been wondering whether MIDI had short-circuited the need for acoustic analysis. But you would say there is still too much happening during the note to trust to sensors?

CC: Exactly. Any instrument can be a polyphonic source, to some degree. Special sensor mechanisms are problematic. You can't impose them on a good player. So we have to have this problem solved acoustically. There are too many applications that depend on it.

On the way to developing a fully competent system, we will be able to explain a lot about how we hear music. There are so many questions right now. It's a fun problem. This is a new type of musical analysis (as opposed to the traditional one where we sit in theory class and pick apart music a la Rameau), looking at the raw fabric of the sound itself. This may allow us to look from a different angle at the higher levels of meaning that people like to talk about. There would be a continuum of analysis. Maybe we could back out from a single waveform all the way to the structural levels of a composition. These kinds of ideas are going to be evolving with the things we find in the next five or ten years.

There is this question of the lifespan of music. There is music that lives on paper, which has analyzable material, as opposed to the music that lives in the air while a piece is being performed. There is an incredible amount of nuance, inflection, and feeling put into it on its way to being realized by the performer. That kind of material, which is added by the performer, is an important part of



music. It's something that was always assumed was going to be added in. For me, a lot of the personal musical moments come from those additions. Right now, we can't see a lot of this data, so we can't form theories about it. Soon we'll be able to get a handle on exactly what they are, and begin really talking about them.

I guess the one other thing to add to all this is the fact that now you can make music by taking advantage of accidents in a way that you never used to be able to. The kind of accidents that come from playing around with sound at the surgical or sculptural level, or playing around with masses of sounds, where you are controlling clouds of notes with algorithms...those kinds of accidents are just wonderful. You find things...you stumble on to them...that have nothing to do with where you were planning to go, but become so important to you that you switch paths.

Like in ceramics. You submit your idea to the kiln, and it comes out looking completely different from what you thought you put in there. Still, it's got a richness that would never have happened without your instigating the whole thing. And it takes you somewhere else. It's fun. It happens in improvisation, to have that kind of serendipity. I like these chances for accidents.

SJ: I'm struck by the amount of cooperation and collaboration that goes on. Everyone is committed to their own interests yet they broadly share techniques and sounds. You don't even have copy protection for files. One might not expect this from a highly competitive research facility.

CC: Each of us is actively engaged in something that is usually not a one person job. Most of what we do depends on teamwork. There is a natural sociology which develops out of that. Progress depends on sharing with others. As Director of the lab, John Chowning is most responsible for this feeling. He has set the example.

You might compare it to the early days of Bell Labs, where you had musicians collaborating with engineers. No one, on their own, would have come up with as many promising developments as they did. If we are committed to anything, it is to this idea that progress depends on a lot of different people with a lot of different backgrounds. It is interdisciplinary by nature. You get interesting, unpredictable results when dissimilar people come in and say they are interested in each other. We have an open policy. We don't limit access to certain backgrounds or qualifications. It just doesn't work. You don't get the results.

There are a lot of people here who are not in the Music department. They may be musicians but they choose to work from within electrical engineering, computer science, psychology, or medicine (speech and hearing). This is important because of the business of the Center. We have to have students who have a lot of engineering chops, as well as psychology, to keep things alive and vital. It's a good place. You probably feel that way about Sequential. It's a good group. Interesting problems. And, you are doing music. Which is the root. It might be time

to leave if I weren't able to do my own music.

SJ: In manufacturing, some are paid to play, but few to compose.

CC: It's business first, here too. But we recognize the necessities and opportunities for doing music.

#### BIBLIOGRAPHY

Computer Music Journal. Cambridge, MA: MIT Press.  
Subscribe now even if you can't understand the articles.  
You will need them later. Besides, its MIDI coverage is growing.

Hal Chamberlin: Musical Applications of Microprocessors,  
2nd ed. Hasbrouck Heights, New Jersey: Hayden, 1985.  
Filled with solid, practical info. A must.

John Chowning and David Bristow: FM Theory and  
Applications. Tokyo: Yamaha Foundation, 1986.

Walter B. Hewlett and Eleanor Selfridge-Field: Directory  
of Computer Assisted Research in Musicology. Menlo  
Park: Center for Computer Assisted Research in the  
Humanities, 1987.

John R. Pierce: The Science of Musical Sound. NY:  
Scientific American Books, W.H. Freeman and Co., 1983.  
A lucid introduction to acoustics. Highly recommended.

Curtis Roads, Ed.: Composers and the Computer. Los  
Altos: William Kaufman, 1985. Interviews with computer  
composers, including Chowning.

Curtis Roads and Jown Strawn, Eds.: Foundations of  
Computer Music. Cambridge, MA: MIT Press, 1985.  
(Paperback, 1987). The most important articles from  
Computer Music Journal's first three volumes. Includes  
Chowning's 1973 FM article.

John Strawn, Ed: Digital Audio Engineering, An  
Anthology. Los Altos: William Kaufmann, 1985. Includes  
Peter Samson's discussion of CCRMA's Systems Concepts  
Digital Synthesizer (the "Sambox").

John Strawn, Ed: Digital Audio Signal Processing, An  
Anothology. Los Altos: William Kaufmann, 1985. Highly  
technical.

#### DISCOGRAPHY

##### Cassettes:

The Science of Musical Sound. Works by Pierce (1959,  
1966), Lewin (1961), Mathews (1961), Moorer (1978),  
Risset (1979), Morrill (1979), and Chowning (1981). \$8.95

Computer Music from CCRMA, Volume 1. Works by  
Chowning (1972), Moorer (1978), Schottstaedt (1978,  
1982), Schloss (1980), Chafe (1981), Berger (1981), and  
Jaffe (1982). \$8.95

Computer Music from CCRMA, Volume 2. 1985 works by  
Marin, Wolman, Nunez, Malouf, and Krupowicz. \$8.95

CD:

Dinosaur Music. Works by Jaffe (1982), Schottstaedt  
(1982, 1985, 1986), and Chafe (1986). \$8.95

LP:

Michael McNabb: Computer Music. 1983. \$6.95

All recordings above are available from CCRMA,  
Department of Music, Stanford University, Stanford, CA  
94305. Add \$1.00 for postage and handling. California  
residents add 7% tax.

CD:

Digital Domain. Elektra

Michael McNabb. Mobile Fidelity.

Forthcoming CD:

John Chowning. Four Works. Wergo.

#### THE AUTHOR

Stanley Jungleib has been with Sequential for eight  
years. He has written more than a hundred manuals and articles  
about synthesis, sampling, sequencing, and MIDI, including  
editing the original MIDI specification itself. His current  
Sequential duties include instrument specifications, user  
interface design, documentation, and marketing communications.

As a player, he generally uses gross amounts of distortion on his  
Rhodes, has been known to haunt the Peninsula's "Friday  
Night Music," and lead the ephemeral electronic  
ensemble, "Pacific Rim."