by Larry Oppenheimer

In "A Byte of Fresh Apple" (January 1990) I presented a view of computers as foundations on which third-party developers build the user environment. The decisions made by the computers' designers both empower and limit the

In no area of application is this more true or important than with sound and music (S&M...love that abbreviation). S&M place extremely stringent demands on a personal computer system, especially with

With that in mind, I examined the motivations and innovations of Apple Computer, the makers of the Macintosh, currently the predominant machine in the U.S. for S&M. This month 1'11 take a look in the door of NeXT, Inc., the latest wunderkind to grab the headlines, and see the 'what and why' (as NeXT's Bob Fraik puts it) of the company's approach to S&M.

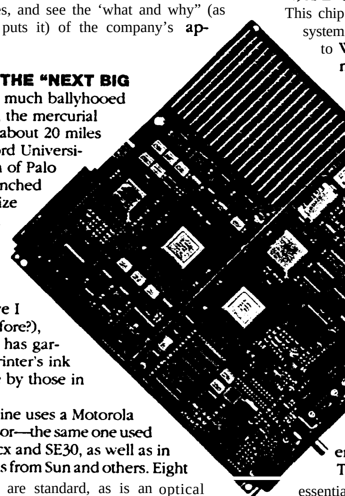

THE START OF THE "NEXT BIG THING" After his much ballyhooed Waterloo at Apple, the mercurial Steve Jobs moved about 20 miles north to the Stanford University-dominated town of Palo Alto, where he launched NeXT, Inc., to realize his personal vision of what a computer should be. Despite considerable delays in its release (where have I heard this story before?), the NeXT machine has garnered gallons of printer's ink and a lot of debate by those in the industry.

The NeXT machine uses a Motorola 68030 microprocessor—the same one used in the Macintosh IIcx and SE30, as well as in several workstations from Sun and others. Eight megabytes of RAM are standard, as is an optical disk drive that holds 220MB removable cartridges. The sound hardware consists of: stereo 16-bit audio playback at a 22/44.1kHz sample rate through a pair of gold-plated RCA connectors, a stereo mini jack and a small monitor speaker, a speechquality (8kHz sample rate) A/D converter; and a Motorola 56001 DSP chip. This chip is the same one used in many digital audio systems, from Digidesign's Sound Accelerator card to WaveFrame's AudioFrame. A fast 12-channel DMA (Direct Memory Access) controller enables the NeXT machine to move the large amounts of data used in digital audio efficiently, thus allowing multiple and real-time processes (like digital recording). Currently, this inventory is the most powerful to be found on a desktop computer system.

Of course, "A computer without software makes a great boat anchor," as the old industry adage goes. Jobs assembled a team of young but seasoned stalwarts from the computer music field (several from nearby CCRMA, Stanford's esteemed computer music center) to create system software the equal of the hardware engine, thus completing the platform.

The new generation of hardware in NeXT essentially puts the power of a mainframe computer on a desktop, perfect for supporting the trend toward more graphically and sonically oriented user interfaces, as well as enables musicians access to tools formerly found only in academic computer music facilities.

(Left): When coupled with an external A/D converter, such as Metaresearch's Digital Ears," the NeXT becomes a digital audio workstation capable of recording, playback and ediing of CD-quality digital audio. (Top): The NeXT computer system. (Above): The motherboard in the NeXT computer, containing both the Motorola 56001 digital signal processing chip and the Motorola 68030 microprocessor.

Composer/programmer David Jaffe is one of the migrants to NeXT's S&M team from CCRMA. "Up until now, computer music has been partitioned into a couple of categories," Jaffe explains. "On the one hand, there's studio computer music [traditional academic computer music studios, which are primarily based around mainframe or minicomputers] in which you have all this control over the sound itself, but the gestural control is limited because of not being able to do things in real time. And you have the MIDI synthesizers with the gestural control, but the timbral fine-tuning is limited because the machines are hard-wired to do one thing, and they only bring out a few buttons. So putting it all in one computer enables you to get at both levels.

"Another division is between sampled sound (or concrete) and synthesis. In the industry, samplers are sort of hybrid things anyway. In a pure sampler you hit a key and it plays this recording, but now when you have a machine that can play Beethoven's *Fifth Symphony* or some sample of long duration, these distinctions break down even further. The way it is now, you have a synthesizer and it does one thing: One does FM, another does sampling. In the case of the NeXT computer, you can design your own algorithms, which used to be something you'd only do in a research facility. Now, experimenting with new algorithms can be an application in itself. We expect an explosion of invention of new algorithms."

Dana Massie, an original member of the S&M group now at E-mu Systems, says, 'Visualization and audio are natural ways to communicate and express things, and this is a platform to help do that. I think that rock 'n' roll is information; it's a form of communication. We expect there will be a lot of very specific applications: voice mail, speech communications, playback of high-quality recordings, generating high-quality music from scratch, processing recorded music, archival recordings of famous people in history...There'll be applications in hypermedia, where you'll be able to have a songbook or a biography of a musician and see pictures of the musicians, or click on the song lyrics and hear individual phrdses of the song played back. Ultimately, the distinctions of the machine as a recorder or processor or mixer or synthesizer are all going to blur."

"Music is something that we all have a big interest in," operating system programmer Gregg Kellogg states, "so perhaps we all focus on that in our designs, but we don't design in restrictions on programming for any applications. The computer is intended for the users out there to do with what they want, and we expect that will change the way we look at this thing in the long run. There's a bigger push toward visualization and being able to have the computer work in more ways than letting you read some text and sort through some things. Three-dimensional color graphics are ways for people to be able to solve problems without having to get into the details of problems themselves, and sound is a natural extension of that kind of idea as well. Sound and music is a lot broader than we really know. Having DSP on the CPU board. for example, means we start getting to a whole new class of speech generation or recognition software that we really couldn't do before."

The NeXT software ream (including the S&M group) is using a two-pronged approach to maximizing the considerable capabilities of the hardware. First, the system software is designed to be as efficient as possible through teleological laziness. i.e., don't do any work that's not absolutely necessary.

Efficient execution is achieved by dynamic loading and scheduling. Given a limited amount of DSP memory, careful management is required to assure that the right sounds are there when the score calls for them. Dynamic loading means instruments that aren't being used at a given time will happily surrender their memory to an instrument that is about to be played and isn't already in memory. When necessary, sounds are shuffled to facilitate collecting fragments of unused memory into the largest chunks possible.

Since the DSP is usually able to outpace the CPU, a large buffer exists between the two to keep the DSP from constantly distracting the CPU with requests for something to do. The CPU can fill the buffer in one fell swoop with messages for the DSP that contain an action to be performed and a time at which it should be done. With this schedule in place, the DSP can compute samples as far ahead of the score as it wants and output each at its indicated time. The system can vary how far ahead of the real time the DSP runs to accommodate different circumstances: With real-time control input (such as MIDI controllers), the DSP must be kept from running so far ahead of real time that it has already computed samples that should be affected by the real-time controls. Real-time control and score events can be integrated in the schedule.

The operating system, Mach, is a key player in making sure that data is where it's needed at the time that it's needed. According to Kellogg, Mach 'is a variant of Unix that provides much better messaging facilities [than Unix] and also control over virtual memory resources, which is a key in our strategy for dealing with a lot of sound files." Virtual memory is a system in which mass storage, such as disks, are treated by the computer as if it was onboard RAM.

"The key is that we don't do any more I/O than we need to,' Kellogg points out. "We delay all the copying of data until we need it." Think of it like this: I want a peanut butter sandwich delivered to my house in San Francisco from my favorite restaurant in Boston, but I need to decide how I want the sandwich cut in half. I could have it sent by overnight express mail, examine it, and express it back to the chef to cut and express back to me, but by then the sandwich would get stale. A better approach. suggests DSP programmer Julius Smith, is to "fax pictures of the peanut butter sandwich and not actually grab it until you're going to eat it." The pictures of the sandwich are messages passed on the network that describe the file in question.

Musically speaking. you could decide to change the arrangement of a song from .4-B-A to A-.4-B. and the structure of the song will have been significantly altered. However. it is not until the song is actually played that any resources (in this case, a musician and instrument) are required to realize these alterations.

Lee Boynton is another key player on the S&M team, who has also done extensive music programming on the Macintosh at IRCAM (the renowned French computer music facility) and MIT. He places great importance on Mach for music and sound work for the reason that, "With enough work, anybody can build a tool that helps you build an application easier. But as far as basic operating system support, that's one of the things you can't really change. Having the basic functionality that Mach provides, [especially] in the virtual memory system, gives you a nice model. It makes a lot of things a lot more reasonable." Besides which, adds Boynton. "I like peanut butter."

The second prong of the approach is providing comprehensive mid-level tools for building applications. This is one place where Jobs' lessons learned from the Macintosh come into play. Although the Mac embodies a high degree of flexibility, it is a bear to program (and a grizzly at that), not because it does not ofier good system support, but because the tools (Managers and Toolbox) are all at the lowest level. There is no inherent mechanism for cresting "primitives" (although there is MacApp, a developer's toolkit available separately, which largely achieves this aim). Primitives can be thought of as system-level macros or batch files, or even extensions to the operating system, which create a level of control just above the lowest by executing a sequence of assembly language routines that accomplish an oft-needed, low-level task.

NeXT's system software is built in layers so engineers can quickly and easily create applications using mid-level tools without mitigating access to the lowest-level nuts and bolts for those situations that demand it. The layers look like this (going roughly from lowest level to highest): DSP software, Mach (the operating system), NeXT Step and the object-oriented toolkits.

The DSP software has two layers of its own. The unit generator layer is based on the model introduced in Music V, a language for programming music created in the 1960s by computer music pioneer Max Mathews and several colleagues at Bell Laboratories [see *Mix*, December 19841. In that scheme, "instruments" are formed by configuring and connecting single-function software modules in an approach reminiscent of analog modular synthesis or. to give a more recent example, Digidesign's *Turbosynth* for the Macintosh. Module functions include oscillators, filters, mixers and basic arithmetic.

As with any modular system, flexibility is the byword. which means that virtually any synthesis algorithm can be implemented: additive, FM, Karplus-Strong and many others. Unlike Music V, however, which chugged through the score computing one sample at a time, the unit generator software on the NeXT machine computes eight samples each pass. The down side is that envelope breakpoints are only computed at the beginning of each pass, but linear interpolation between breakpoints at the sample rate compensates; thus, true sampling-rate envelopes are produced. Also bundled with every NeXT system is Ariel Corporation's BUG-56, the first symbolic debugging software designed specifically for DSP. This package makes use of the NeXT's visual interface to simplify coding for the 56001.

NeXTStep is a set of ***programs-Application Kit, Interface Builder, Objective C, Workspace Manager*** and *Display* PostScript--that sit 'on top" of Mach and allow an application to be built that encompasses DSP software, the Music and Sound Kits, and the user interface.

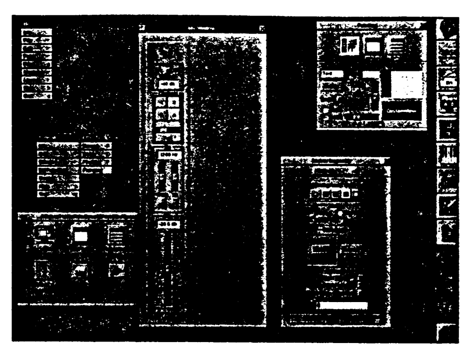

NeXT's Interface Builder being used to design a sectionofa4 channel mixer.

Using the Motorola 56001 digital signal processing chip within the NeXT system, the MonsterScope application is displaying a 440Hz sine wave in Oscilloscope and Spectrum Analyzer formats.

NeXTStep, with the exception of the toolkits, has also been licensed by IBM, presumably for use in its future machines. In fact. future machines are as much the target of NeXTStep as the current NeXT machine, according to Massie. NeXTStep's programs, he asserts, "are based on many years of experience with graphical user interface programming languages. We tried to come up with the next generation [of system software; that might last on hardware platforms for the next ten or 15 years. Ten years is a pretty reasonable lifetime for a major computer [software] architecture." By this measure, the original IBM PC architecture has reached the end of the line, and the Macintosh architecture is' about halfway through.

The NeXT architecture (i.e., NeXTStep), then, has not yet begun its lifetime and will offer significant advantages over previous software approaches to last well into the future—at least that's the fervent hope of NeXT's programmers. Perhaps the feature with the most promise of fulfilling this potential is NeXTStep's portability. An application written on the NeXT machine using NeXTStep can, theoretically, be ported to an IBM machine running NeXTStep, recompiled, and run with all its graphic31 interface features intact. Of course, applications that are heavily dependent on special-purpose hardware, such as sound and music software employing the NeXT machine's DSP chip, can't run without appropriate hardware and the requisite software interfaces to it.

Nonetheless, NeXTStep represents 3 significant step forward in simplifying the task of building a usable application. For example, Interface Builder allows 3 programmer to create a user interface by pulling graphic objects like buttons, sliders and fields off a palette and arranging and resizing them on the screen, and then linking them to the actual application code and defining how they should act. The "look and feel" of an application can be evaluated quickly before the actual code is written, thus shifting programming effort back to the specific code that accomplishes the application's purpose. 'Interface Builder," Boynton says, "is like *HyperCard* squared or cubed."

Objective *C* is an object-oriented implementation of the powerful, popular *C* programming language. With this, programmers fluent in *C* can instantly begin using object-oriented programming techniques. Since Interface Builder is written in Objective C, it is extensible: You can design your own user-interface objects, say. a VU meter, then put them on a palette and use them in your applications.

Sound and Music Kits are libraries of Objective C routines from which sound and music applications will be built on NeXT. 'The distinction is that sound is analogous to bitmaps in graphics or to samples (in sampling instruments), the kind of 'raw form,'" Jaffe explains. "You could say it has 311 the detail but not very much structure. The Sound Kit supports an object interface to the sound capabilities of the computer.' Through Sound Kit routines, basic operations like recording, playback, display and editing are accomplished. Sounds can also be moved on and off the "pasteboard," NeXT's version of the Macintosh clip board. Sound Kit even provides for viewing sounds in several different representations, including waveform display and spectrograms. In other words, the basic sound editing functions we expect from sample editing programs 3re present in the NeXT's system software, and, using Interface Builder's links to the Sound Kit, a basic recording/editing/playback application could be constructed without writing a single line of program code. Further, 3 NeXT sound file can even be instructions for the DSP to synthesize 3 sound, rather than making use of actual samples. And, since the sound file itself is an object that "knows" whether it is made up of samples or code, the application needn't make any distinction when addressing sound objects.

The Music Kit implements objects that embody concepts of musical structure, such as scores, pans and notes, as well 3s orchestration31 elements like instruments, performers and conductors. It also provides access to the NeXT machine's DSP for doing virtually any kind of synthesis, from wavetable, FM and additive to other rarer, more exotic beasts. Fine control of sounds benefits immeasurably from the nature of object-oriented programming, which endows each object with a sort of "cultural* or "genetic" knowledge about its capabilities.

"There's 3 note object that is basically 3 package of parameters, as in Music V, where parameters are things like frequency and amplitude, but an application can design its own parameters that can be 'type of reed' or whatever," Music Kit author Jaffe states. "There's also the idea that the parameters themselves are just information. like a property list, and it's the consumer of that information that assigns the semantics to it. So, for example, to 3 wind instrument 3 brightness parameter might mean to blow harder, and to 3 stringed instrument it means to press harder on the bow, etc. The note object contains information that some interpreter will analyze somehow. That interpreter either responds to that message or not, and, if not. it just ignores it. If he doesn't know what brightness is, he just doesn't pay attention: if he's interested in it, he looks for it and does what he wants with it." Using this approach, a composition (or 3 sound effect) could be orchestrated by deciding on the sound's timbral evolution for 3 given part, then applying that to various instrument sounds until the most appealing choice is found.

The Music Kit, of course, can deal expertly with MIDI vi3 MIDI objects, providing 3 deeper and broader representation of performance that includes MIDI 3s 3 subset. There need be no distinction made between playing sounds generated by the onboard DSP and playing them on MIDI synthesizers. There is even a MIDI file object that allows the NeXT machine to read standard MIDI files as scores.

What the Music Kit doesn't encompass is notation and graphic representations, which, of course, is where Interface Builder comes in again. Clearly, the components of the NeXT machine's system software are intended to be highly integrated with each other.

GREAT. SO WHERE ARE WE IN ALL THIS?

Massie's analysis of computer lifetimes fits with the current state of affairs: As of this writing. Macintosh is in full flower and boasts many applications for music and sound, while NeXT has no off-the-shelf music and sound applications software available to speak of. But NeXT has 3 greater conception of how a computer should deal with music and sound, reflecting developments pioneered by the Mac. And one can confidently gamble that Apple is aware of this and is working on its next generation that will be introduced around the time the NeXT machine attains 3 level of maturity. The onboard capabilities of machines like these make them potent platforms for building 3 music and sound working environment sufficently complete to consider bestowing on it that now overused buzzword: WORKSTATION! ■

When not performing with Phoenyx, Larry Oppenheimer works as a consultant and sound engineer in the San Francisco Bay Area.

Audio Add-Ons for the NeXT Computer

Already a few interesting third-party products have appeared for NeXT. Perhaps the most innovative is the DM-N digital microphone from Ariel Corp. of Highland Park, N.J. It contains a pair of Prim0 capsules (mounted in an X-Y cardioid configuration) that directly feed a Motorola 56ADC sigma-delta analog-to-digital converter. The DM-N's output connects to the NeXT's DSP port. Two mini-jacks on the back of the mic offer the option of substituting any line-level source for either or both of the microphone capsule outputs. The DM-N provides 3 signal to the NeXT that is already in the digital domain and theoretically suffers no degradation from the cable run connecting it to the computer.

Ariel is working on an AES/EBU format of the microphone, as well as plug-in DSP cards for NeXT and other devices that will connect to its DSP port. Ariel's BUG-% symbolic DSP code debugger is already included with every NeXT system.

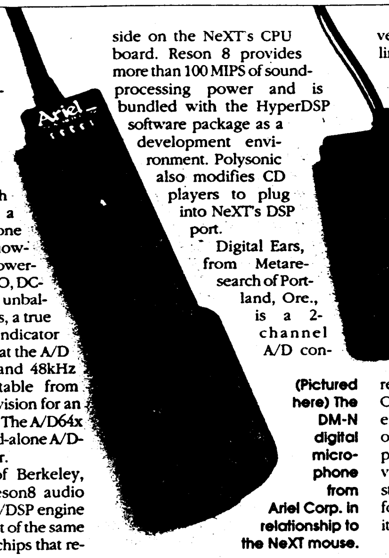

Pasadena, Calif.-based Singular Solutions manufactures a sigma-delta A/D converter, called the A/D 64x, which also plugs into the DSP port. The A/D 64x is loaded with features including: a high-quality microphone preamp, switchable low-cut filter, phantom powering, AES/EBU digital I/O, DC-coupled balanced and unbalanced line-level inputs, a true digital overload indicator (which actually looks at the A/D bitstream), 32, 44.1 and 48kHz sample rates (selectable from NeXT), as well as provision for an external sample clock. The A/D64x is also usable as a stand-alone A/D-to-AES/EBU converter.

Polysonic, Inc., of Berkeley, Calif., makes the Reson8 audio processor, a synthesis/DSP engine built from a set of eight of the same Motorola 56001 DSP chips that reside on the NeXT's CPU board. Reson 8 provides more than 100 MIPS of sound-processing power and is bundled with the HyperDSP software package as a development environment. Polysonic also modifies CD players to plug into NeXT's DSP port.

Digital Ears, from Metaresearch of Portland, Ore., is a 2-channel A/D converter that outputs 44.1kHz, 16-bit linear PCM data into the NeXT DSP port. The device features level controls for its unbalanced RCA inputs and can sample down to DC. Bundled with the Digital Ears is the *SoundWorks* software package, which allows recording multiple takes, editing and selective saving to files with onscreen VU meter animations. Metaresearch also has rewritten parts of NeXT's Sound Object to optimize it for faster operation, and bundles this improved object with the hardware. The company also make Digital Eyes, a video digitizer for NeXT that can store images in several popular file formats and offers a range of digitizing resolutions.

(Pictured here) The DM-N digital microphone from Ariel Corp. in relationship to the NeXT mouse.

—Larry Oppenbeimer