THE NEW WORLD OF DSP

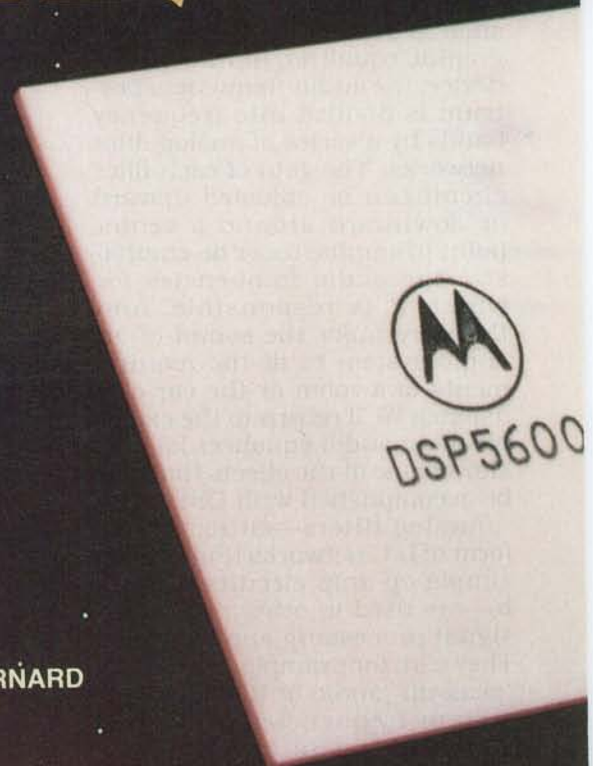

JOSEF BERNARD

WHAT MICROPROCESSORS WERE IN the 1980's, digital signal processing (DSP) will be in the '90s. It is already being used in consumer audio, video, and communications equipment, and will show up in many more products within the next few years. With the ability to design and manufacture fast and accurate analog-to-digital (A/D) and digital-to-analog (D/A) converters, and to process and manipulate digital information at blinding speed, today's technology has made it as cost-effective to manipulate signals in the digital realm as it is to do so in the analog, and the results are better, and more spectacular.

Before we can appreciate some of the devices in which DSP is being, and will be used, we should have an understanding of what it is, and how it works.

Analog signal processing

Signal processing can take many forms. Sometimes it involves changes in signal levels such as the type used in an audio graphic equalizer. In that type of device, the audio-frequency spectrum is divided into frequency bands by a series of analog filter networks. The gain of each filter circuit can be adjusted upward or downward around a center point to emphasize or de-emphasize the audio frequencies for which it is responsible, and therefore tailor the sound of an audio system to fit the requirements of a room or the ear of a listener. We'll return to the example of an audio equalizer later to show some of the effects that can be accomplished with DSP.

Analog filters—either in the form of L-C networks (Fig. 1-*a*) or simple op-amp circuits (Fig. 1-*b*)—are used in other numerous signal processing applications. They can, for example, be used to "peak up" audio or RF signals at certain frequencies or, in the form of high-pass, low-pass, and bandpass or notch filters, to allow the passage of signals of certain frequencies while blocking those of others.

Signal processing can also be used to modify the phase relationships in a complex signal, as is the case in the TINT control found on NTSC TV receivers—although that's not being done digitally, at least not yet. In audio, on a gross scale, phase shifting shows up in the form of "phlanging," a technique used in recording studios to add a rather weird-sounding effect to material. (The technique got its name from the fact that, initially, it was produced by playing two identical tapes and varying their speeds ever so slightly by applying pressure to the flanges of the tape reels. Now, of course, it's done digitally.) And, since the ear is extremely sensitive to the phase relationships of the sounds reaching it and uses them to help establish the location of sound sources, changing those relationships in recorded material can greatly affect the way the sound is perceived (see All About Surround Sound, June 1990 **Radio Electronics**).

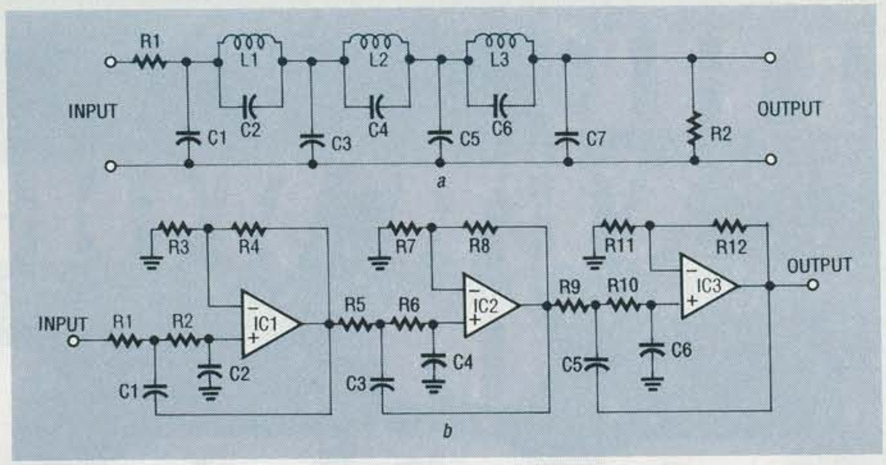

FIG. 1—TWO FILTERS; a passive L-C (*a*) and an active op-amp (*b*). Neither of these types can be as flexible or reliable as one using DSP.

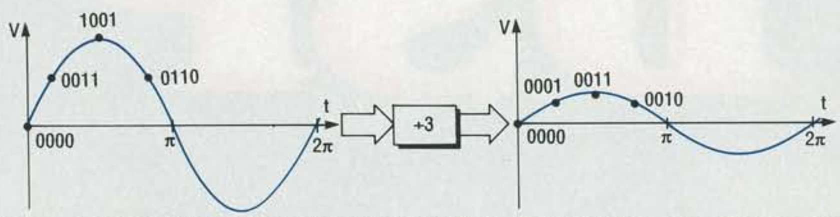

FIG. 2—A HYPOTHETICAL DSP VOLUME CONTROL would divide each signal sample by the same number to reduce it in strength.

Signals can also be summed, or subtracted from one other, to achieve particular results. Both summing and subtractive techniques are used, for instance, in various audio and video noise-reduction schemes.

Until recently, all the signal processing schemes described above, and a number of others not mentioned here, were carried out in the analog domain. Depending on the degree of precision required, the circuits could get very complex and very expensive. Also, a particular circuit could generally serve only a single purpose; if you wanted both frequency-selective processing, for example, and noise reduction, you had to design and build two completely different processing stages. Also adding to its inconvenience was analog signal-processing's dependence on analog components. As analog components heated and cooled, and as they aged, their values drifted and the characteristics of the circuits in which they were used changed. Precision in the analog world can be extremely difficult and costly to come by. Digital signal processing, however, is entirely different.

Digital signal processing

Once you have converted an analog signal to digital form—to a string of binary numbers representing the voltage levels of the signal as it varies over time—you can very easily perform all sorts of operations on those numbers that will affect the signal they represent when they're reconverted to analog form.

Let's take a very simple example. Suppose you wanted a digital volume control, which might, under certain primitive circumstances, be construed as a kind of DSP. To cut the volume of a signal in one third, all you would have to do would be to divide every binary number in its digital representation by three (Fig. 2). The voltages represented by the resulting numbers would be one third their original value, and the amplitude of the reconstructed analog signal would be one third that of what you started with. By changing the divisor, you could vary the amplitude accordingly in either direction.

Taking the process a step further, if you were to multiply all the

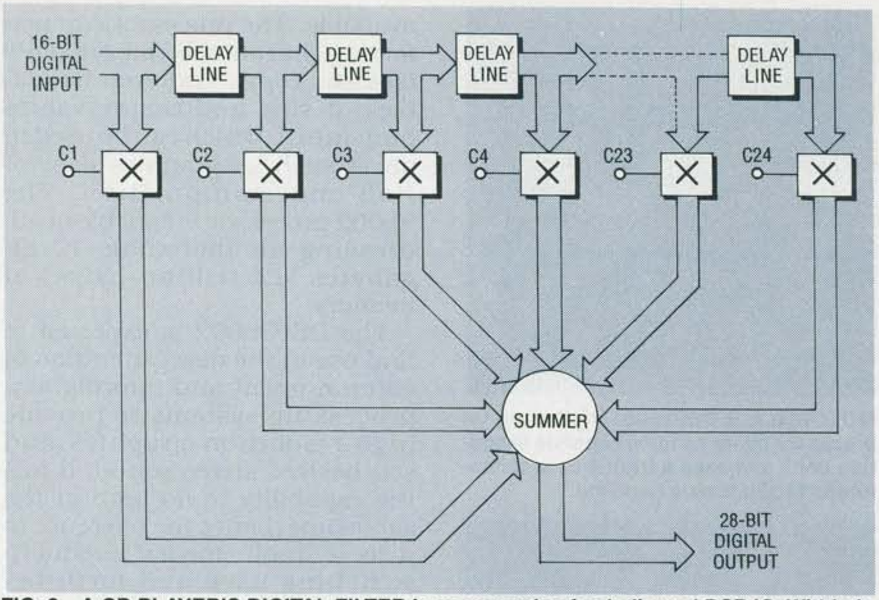

FIG. 3—A CD PLAYER'S DIGITAL FILTER is an example of a dedicated DSP IC. Within it, each sample is multiplied by a fixed coefficient; this is the basis for oversampling.

numbers by one (the equivalent of a unity-gain amplifier) *except* for those representing frequencies between X and Y, which it would divide by, say, 8, you would have a notch filter. That particular filter would, admittedly, have extremely steep sides, but its slope could be modified by using an expression of some complexity to determine the divisor at each point.

As another example of digital signal processing, consider a signal that's stored briefly in RAM as it passes through a system. By reading out that signal a couple of milliseconds after it's been read in, or by shifting it slowly through the RAM's addresses, and then adding it back at a lower level (smaller numbers) than the original as *it* was read out, you'd develop a reverberation effect. Note, by the way, that in that application there are two DSP processes going on at once; time delay and level control.

Digital signal processing of the sort we've just described has been with us for a few years, at least in simple form. For example, digital delay lines have long been used in recording studios. Perhaps the most sophisticated "non-DSP" DSP circuitry is that found in the oversampling digital filters used in CD players (Fig. 3), where every binary number passing through is multiplied by a fixed coefficient. The problem with such DSP devices to date—both single-chip ones and the ones requiring an entire boxful of components—is that they have been dedicated to a single-purpose. They've not been very flexible, which has limited their usefulness. Speed, too, has frequently been a limiting factor. The DSP processes used to enhance satellite photographs, for example, do not take place in real time; it is sometimes weeks before the results are available. (Although much of that delay is no doubt due to the bureaucratic process and long periods of "standing on line.")

Even so, as you've read this description of the principles of DSP, you may have said to yourself, "Hey, I'll bet I could teach my *computer* to do that! Then I could do *anything* I wanted!" And you could, but there would be a hitch. Even today's '386 and '486 desktop systems operating at 30 MHz or more are not fast enough to keep up with the heavy computational overhead demanded by good DSP. DSP, as the term is generally used today, requires that information appear at the output at the same rate it is supplied to the input. Today's small computers would need a lot of help to meet that criterion.

50 megaflops

Fortunately, that sort of help is available. Just as numeric coprocessors, such as the 80387, have lifted a lot of the number-crunching burden from their associated microprocessors (and speeded things up enormously in the process) there are now special number-crunching processors for DSP. Those processors make DSP a real-time process—the modified signal comes out as quickly as it goes in. The difference between the "old" DSP and the new is rather like that between taking your pictures down to the drugstore to be developed and picking them up a week later, and owning a Polaroid 60-second camera.

What makes real-time DSP possible is a new class of IC's from companies such as Intel, Motorola, and Texas Instruments, not to mention a number

MIPS, MOPS, FLOPS

Clock speed is not necessarily an accurate indicator of how rapidly or efficiently a device performs. One processor running at, say, 20 MHz, may not have the throughput of another operating at only 12 MHz. The difference is largely that of the device's internal architecture, instruction set, and other built-in programming. For that reason, performance is often more accurately measured in terms of the number of instructions or math operations of which a device is capable in a given period of time.

The term "MIPS" stands for million instructions per second, and refers to the number of commands that can be executed in that time. A 27-MHz Motorola DSP56001 runs at the rate of about 13.5 MIPS, a 33.3-MHz DSP96002 at 16.65 MIPS.

"MOPS," which stands for million operations per second, is a more accurate measure of a processor's abilities, since a single instruction can be responsible for the simultaneous execution of several (six in the case of the 56001, up to ten for the 96002) operations at once. Such operations can include those for math (add, multiply, subtract), for moving data internally and to and from memory, for carrying out program instructions, and so forth. The 56001 can perform at the rate of 81 MOPS and the 96002, 165 MOPS.

Finally, "megaFLOPS," or MFLOPS, stands for millions-of-FLOPS, or millions of floating-point math operations per second. Floating-point math, which uses exponential notation, is used extensively in the complex calculations required by such applications as 3-D graphics and image processing. Fixed-point math, which is like the "integer math" used in early Apple computers, is much simpler, and cheaper, than floating point to implement and perform and is what the 56001 uses. It is well suited, though, to audio processing where the calculations are not as involved as they can be for graphics.

of offshore manufacturers. Just as the architecture and instruction sets of math coprocessors are created with a specific purpose in mind, these special purpose devices are tailored to the high-speed processing of the digital equivalents of audible and visible information. While each manufacturer has his own idea of what a DSP device should do, and how it should do it, the general principles are the same—real-time manipulation of digital data representing analog phenomena. We'll look at two DSP IC's from Motorola representing the group.

Motorola's DSP56001 is a "general purpose" fixed-point-math DSP IC that's found applications in a number of different types of devices. For example, it is an integral part of Steve Jobs' NeXT computer, that literal "black box," serving to provide on-board data communications (modem and fax) and sound synthesis for such purposes as voice mail, voice-interactive programs, and high-fidelity, CD-quality audio. The 56001 is also incorporated in Cincinnati Microwave's Escort radar detector where it differentiates between radar signals and other, unwanted, types of noise. At a price of $56, even in single-unit evaluation quantities, and maybe less by the time you read this, the DSP56001 is affordable enough to show up in a number of mass-produced general-market devices.

Speed is essential to real-time signal processing, and the DSP56001's specifications demonstrate how it performs in that area. For instance, the processor runs at a speed of 27 MHz, with an instruction-cycle time of 74.1 nanoseconds (0.0000000741 seconds). In the time it takes to execute one of those cycles, a beam of light would travel about 73 feet!

Other 56001 specs include
- Word length of 24 bits, providing a 144-dB dynamic range.
- Capability to execute at 94.5 million instructions per second (MIPS).
- Three complete and independent execution units capable of operating simultaneously and in parallel.
- Triple-bus Harvard architecture. (Harvard architecture, used in some RISC processors, involves two separate buses; the 65001 has three.)
- The ability to perform six separate operations simultaneously.

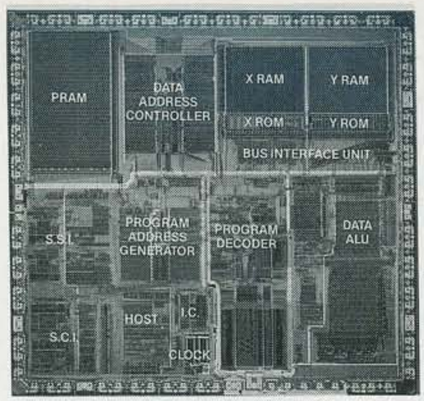

MOTOROLA'S DSP56001 digital signal processor contains three separate execution units and uses a triple-bus architecture to facilitate data handling.

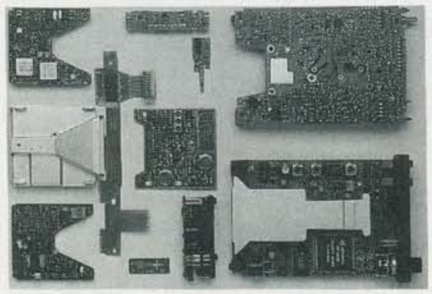

THE DSP56001 at the heart of Cincinnati Microwave's *Escort* radar detector increases its sensitivity by differentiating between radar signals and other types of microwave "noise."

The physical and electrical specifications of this Motorola DSP device are impressive, both in terms of large and small numbers as well. The DSP56001
- Comes in an 88-pin package.
- Operates from a single 5-volt supply.
- Has five 5-volt and seven ground pins to ensure even power distribution and glitch-free operation.
- Consumes less than half a watt of power.

An even more powerful, although somewhat more specialized, DSP device is Motorola's DSP96002 "Media Engine." This 32-bit device, whose internal accumulators can store numbers 96 bits long, comes in a 223-pin PGA (pin grid array) package 1.845 inches square, and uses a one-micron architecture to accommodate some 850,000 transistors. It operates at a clock speed of 33.3 MHz, although a "slow" 27-MHz version is also available. The processor can perform at the rate of 50 megaFLOP's (see box copy), and even has tables of sine and cosine values built into it, which can be used in areas such as graphics generation and manipulation. The 96002 processor is capable of addressing an incredible 12 gigabytes (12 trillion bytes) of memory.

The DSP96002 is expected to find use in the new generation of entertainment and information-processing systems to provide high-resolution graphics and synthesized stereo sound. It has the capability to do both at the same time (hence the reference to it as a "multi media" product), switching back and forth between the two tasks so quickly that no interruption is apparent. In medical, and other, imaging technologies DSP will prove itself invaluable in enhancing and manipulating visual data. The Motorola IC is also expected to find application in color laser printers where it will convert page-description-language commands into the fonts and graphics that appear as output, and as controllers in huge high-end computer disk drives that require constant compensation for the effects of thermal expansion and vibration. The science of robotics, too, will benefit from the ability of a processor such as the 96002 to perform powerful floating-point calculations in real time. For large and complex tasks, several 96002's can easily be configured to operate in tandem and divide the work into more manageable slices, apportioning it among them.

Motorola has plans to introduce an entire family of DSP products. One of the first is the 56ADC single-chip analog-to-digital converter. It can process signals at the rate of 6.4 *million* samples per second (Ms/s) (CD's, in comparison, use a sampling rate of 44.1 *thousand* samples per second), eliminating the need for complex sample-and-hold circuitry. Also in the works is a "sawed off" 16-bit DSP device, as well as a 40-MHz version of the 32-bit 96002.

Some real DSP products

A very good example of some of the ways DSP will be showing up in consumer electronics equipment is Sony's STR-D2010 stereo receiver. In it, many analog functions have been replaced by their digital equivalents, and in implementing those Sony has added an extra degree of flexibility to the features available to the user. A proprietary DSP IC, coupled with 16-bit A/D and 18-bit D/A circuitry and an 8× oversampling digital filter (which, as we've pointed out, in itself provides a form of DSP) is firmware-programmed for a number of useful operations.

The receiver has no treble or bass controls. Instead, it contains a parametric equalizer to tailor frequency response. Parametric equalizers used to be pretty tricky to design and use. They differ from "ordinary" graphic-style equalizers in that the frequency bands on which they operate, and their response curves within those bands, are adjustable to suit the needs of the user. The device's parameters of operation can be changed by the user. With its DSP IC (and 256K of on-chip RAM) the STR-D2010 allows you to define three separate frequency bands (a vacuum-fluorescent display allows you see the response curves, and also functions as a spectrum-monitor display), each with its own degree of boost or cut, and with one of several slopes. There's no "loudness" control either. Instead, a digital signal-compression technique is used to compensate for the way the ear perceives sound at low volume levels.

The Sony unit also has surround-sound capabilities, and uses DSP techniques to provide the various signal delays used to manufacture ambience through artificial echo and reverberation. There's Dolby Surround processing too, using digital techniques to extract the matrixed surround information from the left and right-channel audio signals.

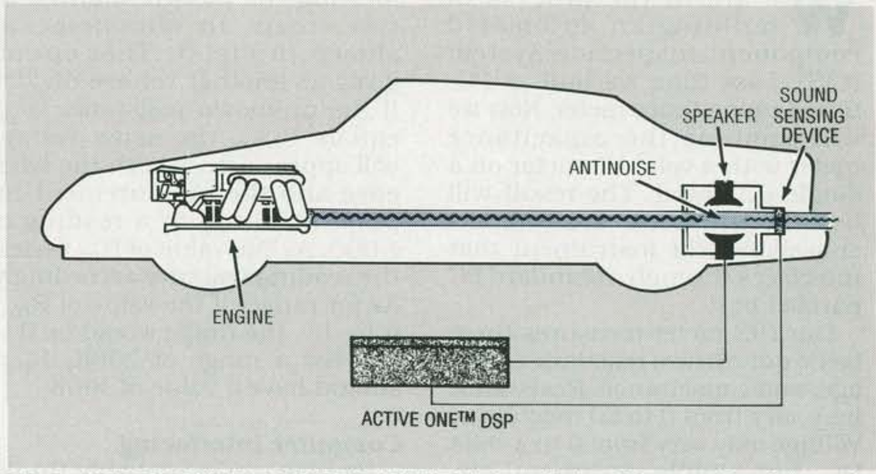

FIG. 5—USING DSP NOISE-CANCELLATION, this "stealth" muffler not only quiets engine noise, but improves performance up to fifteen percent by doing away with conventional baffle systems.

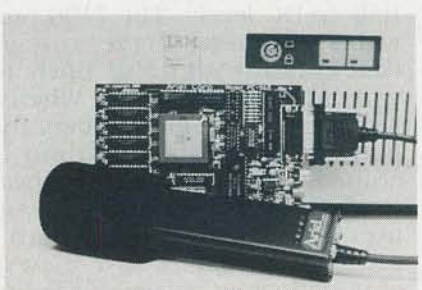

FIG. 4—A DIGITAL MICROPHONE, known as the DM-N from Ariel Corporation uses two Motorola A/D converters to provide two channels of digital input to the NeXT computer system. It operates at the rate of 6.5 Ms/s, eliminating the need for sample-and-hold and anti-aliasing-filter functions.

An autosound receiver from Eclipse includes DSP technology to provide ambience and other effects. Video, too, can benefit from DSP, although the term is generally not yet applied to the processes being used. Video noise reduction, for example, can be accomplished by digitizing the video and comparing successive pixels or adjacent lines. By pixel or line averaging, or even replacing a "bad" pixel with a "good" one, picture quality can be improved. It is even possible to average two or more successive *fields* of video to smooth things out.

Amateur radio is getting into the DSP act, too. Kenwood's top-of-the-line amateur transceiver, the TS-950SD, uses DSP in a number of ways to enhance incoming and outgoing signals. The receiver section, for instance, includes such DSP features as a digital AF filter with user-variable characteristics. The transmitter uses digital techniques for speech compression that increase average output power while keeping peak power output the same. It is also possible for the waveshape of the signal to be manipulated using DSP techniques to increase intelligibility.

IC SOURCES

The Motorola DSP96002 digital signal processing IC currently sells for about $750 in single-unit evaluation quantities, $650 for the "slow" 27-MHz version. More information on this and the company's other DSP products is available from:

Motorola Microprocessor Products Group
DSP Marketing
6501 William Cannon Drive West
Austin, TX 78735-8598

One supplier of complete sets of pre-programmed DSP IC's is:

The DSP Group, Inc.
1900 Powell Street, Suite 1120
Emeryville, CA 94608

The manufacturer of the DM-N digital microphone and MM-96 multimedia board for MS-DOS computers, can be reached at:

Ariel Corporation
433 River Road
Highland Park, NJ 08904

Finally, the company that has developed the DSP "stealth" muffler is:

Active Noise & Vibration Technology
3811 E. Wier Avenue
Phoenix, AZ 85040

On the cutting edge

Besides the applications we've already mentioned, DSP is now being used commercially in such devices as cellular telephones for compansion; in phone-answering and cordless-phone equipment for speech digitization, synthesis, and storage; in transcription and dictation units for variable-speed playback; and in facsimile machines and other de-

continued on page 53

NEW WORLD OF DSP

continued from page 47

vices with modems for data compression. At least one company specializes in the preparation of entire sets of DSP IC's to meet particular signal-processing requirements.

Two Motorola 56ADC's are at the heart of Ariel Corp.'s DM-N digital microphone (Fig. 4), which is designed for direct connection to the NeXT computer system, and the same company already has designed an IBM/AT-compatible plug-in board that contains a pair of DSP96002's. This board, the MM-96, can provide a microcomputer with mainframe-class multimedia performance for scientific, industrial or artistic applications.

And a company in Phoenix is working on a DSP-controlled muffler that is said to improve automobile performance by between eight and fifteen percent! Using electronic noise cancellation techniques, this "stealth" muffler (Fig. 5) does away with baffles and other obstructions, allowing exhaust gases, but not the noise that usually accompanies them, to pass straight through. A fringe benefit of a DSP muffler would allow you not only to silence your engine, but perhaps also to tailor the muffler's output to make your Chevy Nova sound like a Ferrari. Work is also being done on silencers for such notorious noisemakers as helicopters. Shades of Blue Thunder!

The DSP technology and applications we are seeing now are just the beginning of what will prove to be a significant era in electronics. Consumer products that include DSP circuits are soon going to be popping up like flowers blossoming after a desert rain. Some applications are ripe for DSP now, but some of the devices that we will soon see will perform functions of which we have not yet even conceived. Some of the uses to which DSP will be put will be ingenious, and some of them absurd. And a few of the applications will have a significant and long-term effect on the way we conceive of and use electronics. **R-E**