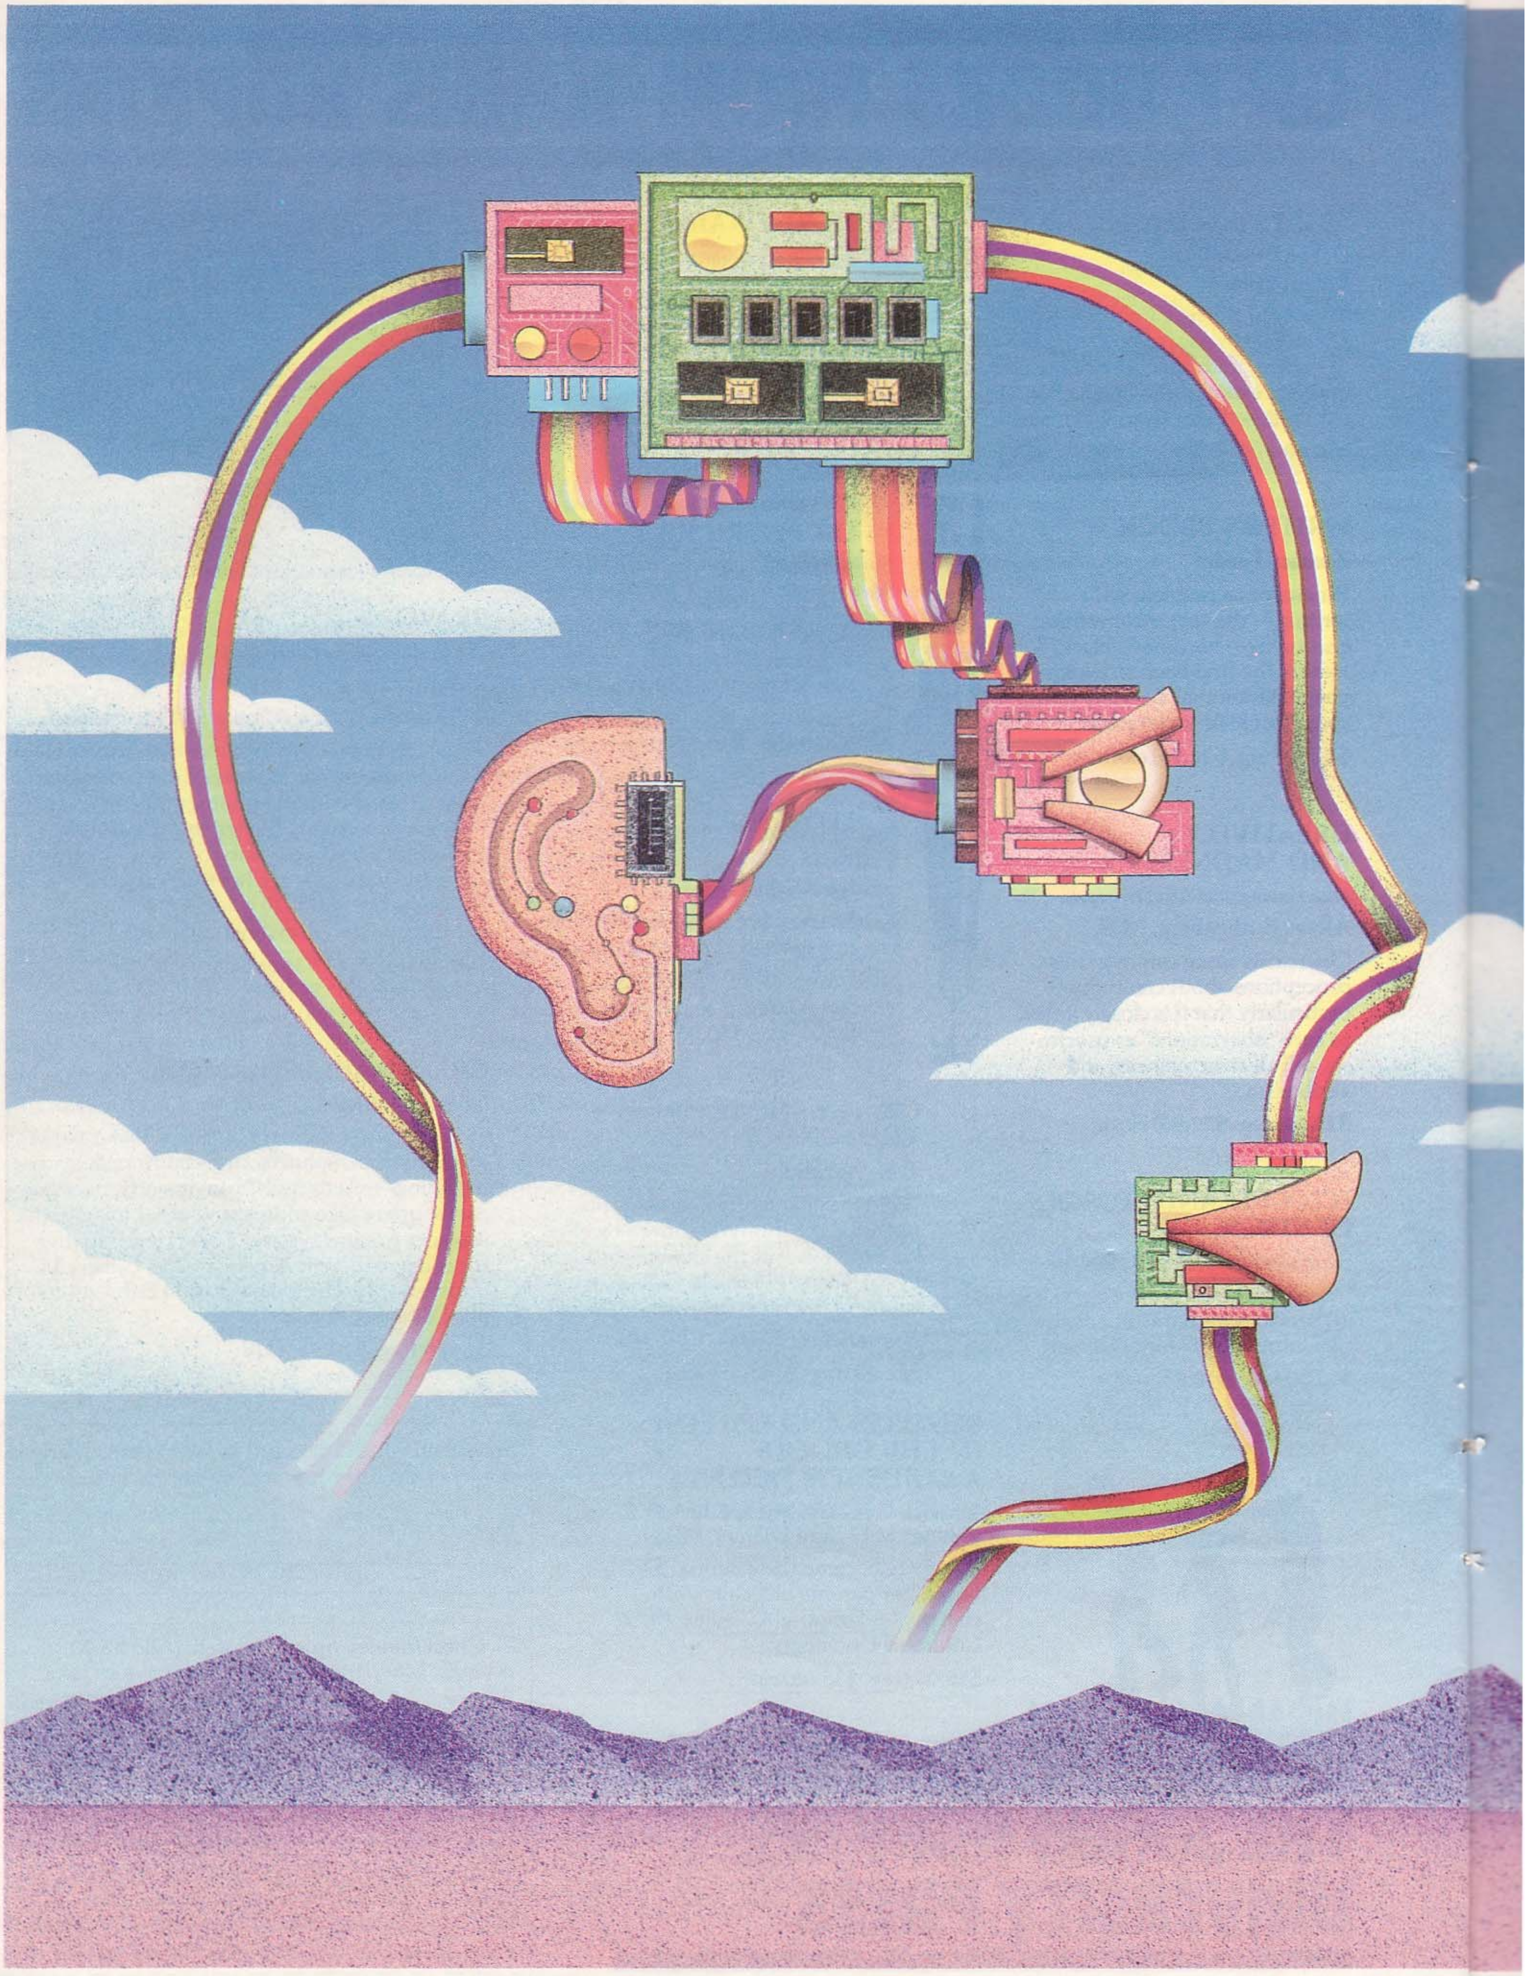

ILLUSTRATION: JIM OWENS

A Chip You Can Talk To

... and it will obey, right away. That is the goal. Oh, and it ought to be small enough to fit in your phone.

By Rachel Nowak

"I DO NOT LIKE TO DRIVE A CAR. I THINK IT WOULD BE nice to get into my car and say, 'Take me to Washington, to the Kennedy Center.' And I could just sit there and read something," says Andreas Andreou, PhD '86, assistant professor in the Department of Electrical and Computer Engineering (ECE).

"I hate phones. I don't remember numbers. Why can't I say, 'Dial Silvio in California'? Why does a telephone have to have a dial, or buttons?" Andreou swoops a telephone from his desk and hammers on the buttons. "It already has a microphone, something to listen with!

So far, large-scale computing is digital. But at Hopkins, one group is exploring another approach: analog computing. Already, they are making chips that see, hear, and learn.

"And this stupid thing!" Andreou grabs my tape recorder. Why couldn't a computer directly transcribe from audiotape into typed copy, he wants to know.

A thickset, robustly healthy Greek Cypriot, Andreou is exasperated by what he sees as the major flaw in traditional computer technology: its inability to "diffuse through society and help people." According to Andreou, computers grow ever more powerful, but remain instruments of the specialist—the number-crunching scientist or the symbol-shuffling word processor operator. "I would like to see that change," he says, hands moving, words tumbling with excitement. "I would like to see people having access to things that

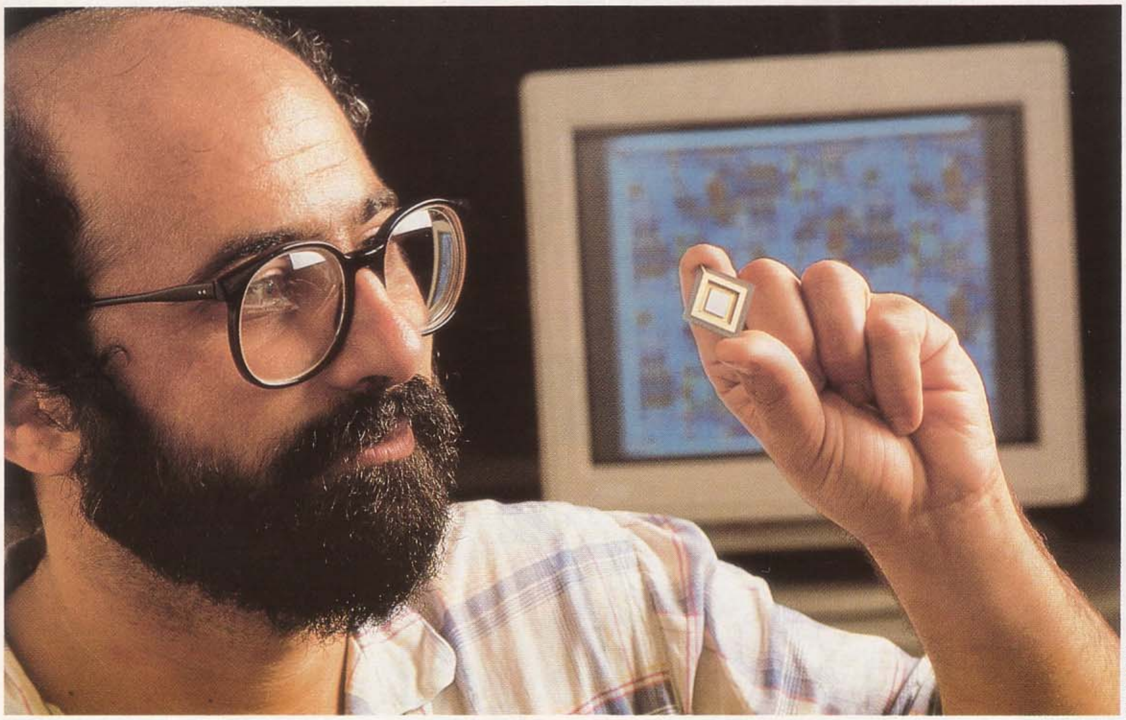

BILL McALLEN

Apostle for analog: Andreas Andreou holds an analog chip that "hears" and sorts one voice from another, instantaneously. Someday, maybe, its distant descendant will be in your phone. Andreou says digital computing is too bulky, too slow, for many devices of the future.

will make their lives easier."

For computers to become useful tools for the masses, they must be redesigned, so they communicate not via the keyboard, but directly with their operators and their surroundings. They must obey the spoken word, recognize faces and objects, even read handwriting, Andreou reasons. To do that, they need to be given sight and hearing, possibly also touch and smell.

Robert Jenkins, group supervisor of the computer science and technology group at the Johns Hopkins Applied Physics Laboratory (APL), shares Andreou's dreams. In the fall of 1985, they met: Jenkins a visiting professor on loan from APL to the ECE department, and Andreou a graduate student in that department.

"There was an instant symbiosis, a meeting of minds," says the slow-spoken Jenkins. With Moise Goldstein, Edward J. Schaefer Professor of Electrical Engineering and an authority on auditory processing, Andreou and Jenkins embarked on an exploration of analog computing for solving problems in sensory perception.

Now they orchestrate the projects of a team of graduate students from the ECE department, who beaver away in a laboratory on the second floor in Barton Hall. The students are striving to produce perceptive computers—perceptive meaning they can translate streams of data, like sound or light, directly into meaningful information.

THE FIRST STAGE IN the team's quest for computer perception is to develop a machine that will respond to verbal commands uttered in the typical noisy environment in which we live and work. The machine would ignore footsteps and the conversation by the coffee pot, but respond to "Dial Dr. Aleksandra Pavasovic, in Belgrade," or "Make a transcript of our conference call." And to be of practical use the speech recognizer machine must operate in real time—that is, must process sound waves immediately, at the same time they are received—as well as be small and energy-efficient.

Most computers in use today are digital and serial. To get an answer, they take a digital input, a string of zeros and ones represented by two different voltages, and then use a mathematical recipe called an algorithm to plod through an orderly sequence of logical manipulations. The problem is that speech, to a computer, is not digital and serial; it is a jumbled mixture of electrical signals coming from a microphone. The signals are analog. That is, the voltage goes up and down, varying through all degrees between on and off. Thus, conventional speech recognizers must first convert this analog signal into a digital string of ones and zeros. That takes time.

Next, the computer goes to its central memory, to compare the string with sequences representing certain words or word parts. If the word has several different meanings, like "ball," the computer has to grind through yet another sequence of logical steps, examining associations with other words in the sentence—"park" and "bat" or "music" and "gown"—to determine which definition fits all the associations. As the sentence gets longer, the number of combinations increases exponentially, a phenomenon known as "combinatorial explosion." The computer, having to scan more and more data, takes longer and longer to search for the answer; it may become completely disabled. Or as Jenkins puts it, "It if's and then's itself to death."

The process is clumsy, time-consuming, requires megabytes of memory, and makes conventional speech recognizers notoriously inaccurate unless the number of possible words is severely restricted, say to a few hundred. Programs that can analyze unrestricted speech are in the research pipeline, but so far they are incredibly slow, taking many hours to process a single segment of speech, says Andreou. The conventional approach to this problem is to build "super" computers—ever faster, but still digital.

Andreou and Jenkins stand among the minority in the field who believe that the conventional "brute force" approach will never produce practical machine perception. Practical applications require that all processing be carried out in real time, they point out. (That is, simultaneously.) And they believe that digital computers—no matter how fast—are inherently handicapped by their need to conduct computations in lock-step order. Andreou says, "We know [machine perception] is not doable with present-day computers. We know it won't be doable with the next generation of supercomputers. So we have to look at different ways of doing things. Radically different ways."

They turned to nature for their inspiration. For while conventional computers may outshine humans at tasks such as logic and math, humans do far better on skills needed in perception—making generalizations, for instance, or finding associations between things previously unrelated.

That's because humans, like all animals, possess sophisticated machinery for monitoring the environment: biological processors that operate in real time, sorting through reams of fuzzy, noisy, environmental data. That's how we make sense of the thousands of sights, sounds, smells, and textures with which we are bombarded. Nature's system is effective, robust, and energy-efficient—a perfect model to follow, say Andreou and Jenkins.

Consider:

The nervous system does not resort to converting analog signals into ones and zeros. Real world data—sound waves, light waves, and so on—are broadcast in continually varying analog form. The ear and the eye pick up these signals and convert them to analog electrical signals, which are then processed again. The ear and eye sift out the important information and transmit it to the brain in the form of nerve cell firing rates.

The nervous system does not require complete sets of accurate data in order to process information correctly. Instead, it uses intelligent guesses based on generalizations, context, and best fits. In this way it avoids the problem of combinatorial explosion. For the brain, the longer the sentence and the greater the number of associations, the more quickly it knows the correct meaning of the word "ball."

The brain, unlike the digital computer, stores information primarily in analog form. Information is stored when complex electrochemical signals change the strength of synapses between the neurons; that is, the signals actually and physically change the way these specific neurons will transmit signals in the future. In this way, the brain does not require programming, but simply learns by experience as it responds to the environment.

The physical layout of the brain ensures that it is robust and long-lasting. Brain cells, or neurons, are not serial, linked in a chain, but are interconnected in a dense web. And rather than each neuron having one specific task, many neurons fire in concert, working together on each job. This cooperation is called distributive, or parallel, processing, and it makes for extremely fast computation. It also means that signals continue to be transmitted and processed even if some neurons cease to function—and a good thing, too, because from a starting total of 200 billion neurons in the human brain, around 50,000 die each day.

Finally, the brain is extraordinarily compact and energy-efficient. The relatively simple data processing done by the tiny eyes of the fly would require a computer as big as three refrigerators running on kilowatts of power, if conventional digital computation was used.

The objective of the Hopkins engineering team is to transfer some of these desirable traits of the nervous system into silicon. To do this, they are designing microchips embedded with analog circuits. As in the brain, analog circuitry allows for speed, small size, and low energy consumption.

The team hopes eventually to assemble the microchips into a miniature, real-time, low-power electronic silicon speech recognizer that will convert spoken words ("speech signals," to Andreou) into the electronic signals that operate a computer. To do that, they must find ways for computers to perform each step of the perception process: *hearing* a command; *separating the commanding voice* from other noises; augmenting sound with *vision*; and finally, transforming words into commands into actions, by use of *associative memory*.

Designing an electronic ear

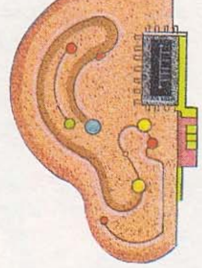

GRADUATE STUDENT Weimin Liu sits amid a mass of electronic apparatus: a screen with three neon green shapes writhing upon it. Boxes with switches and dials. Two fingernail-sized microchips, each centered on a circuit board cluttered with colorful electronic devices like those that spill out the back of old radios. And a tape recorder blasting out pop singer Sinead O'Connor's "It's been seven hours and 15 days...."

Liu is making an electronic ear.

Sound waves entering our ear hit the ear drum, he explains. Then they vibrate along three small bones in the middle ear, through the "oval window," and into the inner ear, a fluid-filled chamber called the cochlea. Vibrations in the cochlear fluid move a membrane, called the basilar membrane, that is stretched across the inside of the cochlea. The movement bends tiny stereocilia on the hair cells that extend from the membrane to nerve fibers of the auditory nerve, generating signals for relay to the brain. It almost seems mechanical, says Liu. "You bend the stereocilia and that generates the signal."

Different clumps of fibers respond to different frequencies, so that different sounds trigger different nerve fibers to produce characteristic electrical firing patterns. In this way the cochlea sorts—or preprocesses—sounds at the same time it transmits them to the brain.

"What is so striking" about the ear, says Liu, "is that each group of fibers is locked to a certain frequency." One major group of fibers responds to the fundamental frequency (or tone), while other distinct groups are sensitive mostly to harmonics. The separated signals convey information vital for speech recognition. For instance, vowels can be identified by their harmonics.

Liu's electronic ear, or speech preprocessor, is intended to imitate this natural system. It responds to sound by reproducing the characteristic firing pattern of the auditory nerve fibers.

Liu is from China, where he worked

If graduate student Weimin Liu succeeds, the result might be a new-generation cochlear implant, one that would "hear" and pass signals to the auditory nerve of a deaf person. Here, Robert Jenkins from APL provides technical advice.

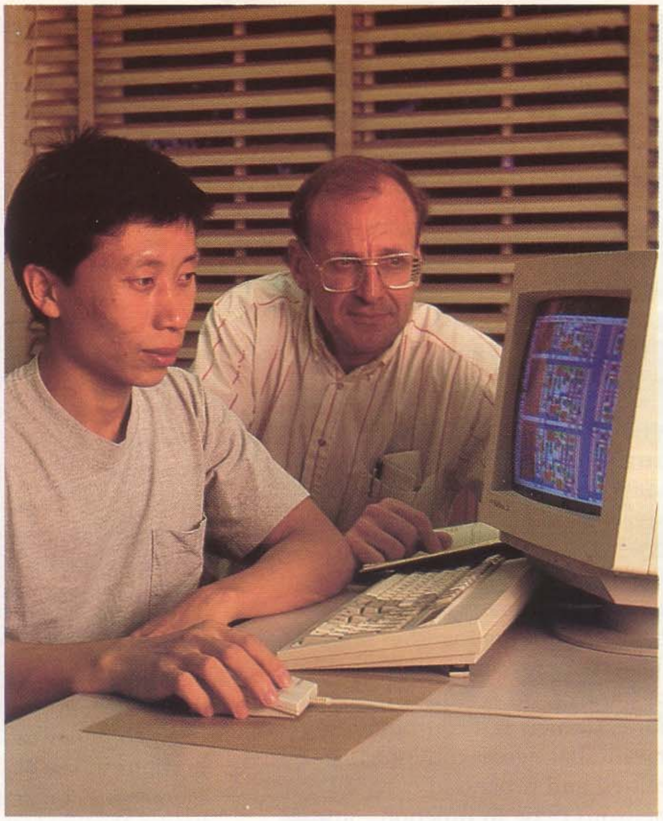

BILL McALLEN

at the Chinese Academy of Sciences in Beijing. He came to the U.S. three years ago to work with Goldstein in modeling auditory processing, and has been working on the electronic ear for about two years. "It may take another two years," he estimates.

On one circuit board sits the basilar membrane chip, a microscopic bank of 30 filters etched into its silicon. Each filter corresponds to a different position along the basilar membrane. (A filter is an electronic device that damps out all but a particular frequency range.)

Each of the 30 filters can connect to an electronic hair cell on the hair cell chip, sitting on the other circuit board. The hair cell chip takes a wide-range signal and compresses it into a narrow range, simulating the interaction of biological hair cells and auditory nerve fibers.

In earlier tests, he passed "da," "ma" and "ba" sounds, synthesized for consistency, through the system. (Newly designed chips, like infants, find these simple syllables "sort of easy" to start with, he says.) In those tests, the signal "locked" into the required frequency, producing a shape similar to the sovereign standard: the firing pattern recorded from auditory nerve fibers of an anesthetized cat "listening" to the same sounds. These preliminary tests were deemed successful.

Today, listening to the music of Sinead O'Connor, Liu is not working. He is playing, to see what his system will do with such complex sounds. He turns on the music, and three green signals dance on the oscilloscope screen. The top signal, leaping and spiking like storm interference on a television, represents the input signal: "Music or speech is a very rich signal," Liu comments.

The filter siphons off its own particular frequency, creating the middle signal, a jaggedy wave, which represents the basilar membrane's output. The signal passes on to the single hair cell on the second chip, and the bottom signal, a smooth, calm wave, represents the final output. Liu smiles.

Eventually, he promises, all 30 filters and all 30 hair cell circuits—or more—will be contained on one chip. But "I want to be sure it works before..." He pauses. "The Big Chip."

"The Big Chip" would be tiny. It would consume little energy (thus needing only a minuscule battery), and it would work in real time to generate a signal close to that of an actual cochlea. If Liu succeeds, the result might be a new-generation cochlear implant, one that could be connected to the auditory nerve of a deaf person.

The team's ultimate aim, however, is to create hardware that can change sounds, instantaneously, into something a computer can understand and act upon: "Call Silvio in Los Angeles," for instance. In engineering terms, the goal is to preprocess a speech signal into an electronic signal that can go into a system that can recognize and even translate speech.

Separating one voice from the din

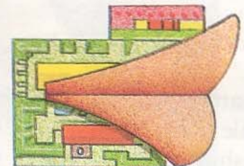

BOMBARDED WITH sounds—the hum of air-conditioners, ringing telephones, droning traffic, cacophonous conversations—a person simply picks out what is relevant to the moment and ignores the rest. But for a computer, sounds coming from more than one source, translated into electrical signals by a microphone, summate like waves traveling through water: they blend. No known method of analysis can separate out—in real time—two unknown signals that have become meshed, says graduate student Marc Cohen.

Behold a microchip, a 2mm × 2mm sliver, thin as a hair. Under the microscope it looks like an aerial view of a glass city silvered by early morning sunlight. Delicate, but strong-looking wires fan out from the edge of the chip, bridging the gap between microcosm and macrocosm.

This is the microchip that can solve the problem—or so Cohen hopes. This chip, a very early prototype, is made to his design. After working at M.I.T. for five years, he came to Hopkins two years ago to do graduate work in the Department of Biomedical Engineering. There he worked on a project recording electrical signals from monkey brains, but realized that to make sense of his reams of data, he would need a way to separate, in real time, the electrical signals from individual brain cells. After falling into conversation with Andreou, he joined the Sensory Communication group a few weeks later.

On this particular day, because Cohen is visiting his family in Johannesburg, South Africa, Andreou acts as guide, playing a prerecorded tape of Cohen's chip in action. On the tape, a voice drowned by a hubbub of music is fed into the chip. It emerges transformed into one clear voice with a party going on in the background—

like a movie camera zooming in on the main character. Another prerecorded tape: Two voices, babbling one over the other, are separated out into Andreou speaking Greek and Andreou speaking English. The speech separator can be adapted to separate out any number of voices, says Andreou.

"The chip looks at the signals coming in: each little sequence, in each little interval of time. It finds the differences and the similarities between them. It carves out little pieces from the pattern in time until both channels have nothing in common," explains Andreou.

The circuit layout used in the voice separator is a type of neural network. (Some engineers prefer the terms "parallel distributive processing" or "neurodynamic computing.") "Neural network" is a catch-all phrase generally used to describe computer systems that, in a crude caricature of the brain, consist of interconnected webs of simple electronic processing units called "neurons." Cohen's voice separator chip, effective though it is, consists of only two neurons and two synapses. ("Just think what it can do when it's got more than two!" says Andreou.)

A neural net is just an architecture, and it can be implemented in a variety of ways, Cohen explains later. For instance, you can simulate a neural net on a traditional sequential digital computer, if you want to. But Cohen sees no point in it, for this project. "The whole idea of doing something on a neural net is that you can perform parallel processing, and that can't be simulated." So he has actually built his neural net into the hardware, in analog form.

Neural networks may or may not work "like" the brain, but are similar in that they need no programming and no central memory. Instead, they "learn" by being exposed to data, and information is stored within the net.

How a network learns: In one process, called supervised learning, the network is exposed to an input signal on one side, and the desired output signal on the other side. Since the neurons are interconnected, the two signals resonate back and forth. As they resonate, the currents and voltages shift on all neurons throughout the net. Special learning algorithms change "weights" (charges on capacitors) to alter the flow of current, or the voltage, between neurons, favoring the passage of input signal to the desired output.

The voice separator uses a more

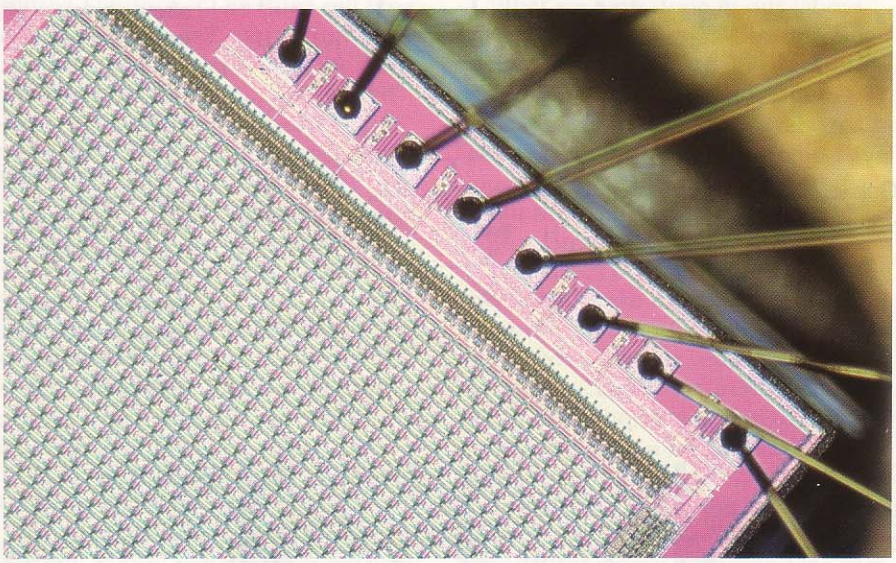

MIKE MANDELLA

Under the microscope, a microchip looks like an aerial view of a glass city silvered by early morning sunlight. Wires bridge the gap between microcosm and macrocosm.

sophisticated method, called unsupervised learning, in which changes in the input signals themselves modify the weights. In this way the system adapts continuously to changing conditions.

In both processes, the input signal can be envisaged as eroding channels of least resistance through the network. Those channels become its memory.

Mathematically, Cohen's chip remains difficult to describe, if not impossible. But that worries neither Cohen nor his professors. "Just build things, test them, and figure out exactly how they work later," Andreou advises.

At present, the voice separator is far from perfect. For instance, it cannot yet deal with time delays caused when speakers stand at different distances from the microphone. Cohen is working on that.

Developing a silicon retina

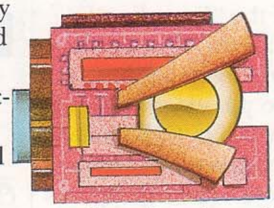

BUT IMAGINE Cohen's chip in a factory, or an airplane cockpit, or even a living room with the television on. In noisy places like that, many sound signals are blocked out, so that speech recognition based on sound alone becomes extremely difficult. "We needed to figure out a way to improve speech recognition in noise. And when you look at the way humans do it, you realize we use a lot of additional sources of information, like hand gestures. Or we just watch the speaker's lips," explains Andreou. But to read those clues, these new computers need vision.

The eye, like the ear, does not merely record the incoming signal. It preprocesses it, in real time, automatically doing things like adjusting to different light intensities, locating the edge of objects, or detecting motion—all that, before passing the sorted information on to the brain.

Conventional machine vision aspires to no such heights. It simply records an image with a camera and converts the intensity of light at each point into a digital signal for processing in a digital computer. Using this method, a computer needs the entire night to deduce that it is looking at two squares moving in opposite directions.

As a step toward real time processing, Kwabena (Buster) Boahen, then a Hopkins undergraduate from Ghana, fabricated a "silicon retina." This device, though lacking the complex skills of a genuine eye, has mastered the heart of the matter: It can convert analog changes in light intensity into analog voltage changes. In that way it can detect edges, as well as automatically adjusting to different light intensities. "It even falls for optical illusions," says Andreou. (In fact, that's how they test these chips.)

Boahen's silicon retina is a neural net on a microchip inspired by one designed in 1987 by the guru in the field, Carver Mead, who is the Gordon and Betty Moore Professor of Computer Science at the other center of this work, the California Institute of Technology. (Gordon Moore is the chairman and founder of Intel Corporation, and one of the "Founding Fathers" of Silicon Valley.) After getting his master's degree, Boahen left Hopkins this fall to pursue a PhD in Mead's research group in California.

The silicon retina has three layers of electronic neurons. The top layer is "photoreceptors," made from a special sort of electrical junction which produces electrical current in direct proportion to the intensity of the light.

As with biological retinas, every point on the processor gets a different amount of light from each part of the image, and each signal is processed concurrently and in real time. The output, representing the logarithm of the brightness, is in voltages. These pass on to the rest of the net for further processing.

In related work, the Hopkins engineering team has produced several different retinal chip models. One will be used in the telescope at the Sacramento Peak land-based Solar Observatory to compensate for movement in the objects being observed. That chip, just one centimeter square, will replace a refrigerator-sized rack of computers.

Associative memory: the final link

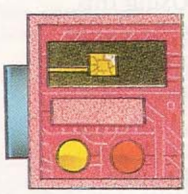

PHILIPPE POULIQUEN, an exuberant first-year graduate student, is the youngest member of the team but by no means the newest. He's been working with Andreou and Boahen since he was a sophomore, developing the last link needed for the speech recognizer—the machine that will transform vowels and consonants into words, words into commands, and commands into actions. The linking device is called an associative memory, and in its present form, it consists of

Making the Silicon Retina

How the chip is designed, and how it imitates nature's "circuitry"

By Andreas Andreou

THROUGH PROGRAMS established by the federal government, university students and researchers have access to state-of-the-art Metal Oxide Semiconductor (MOS) silicon technology. MOS is the most common technology used to implement the microprocessor of a personal computer or other Very Large Scale Integrated (VLSI) systems.

At Hopkins, undergraduate and graduate students can design and fabricate their own analog or digital VLSI circuits through courses offered in the Electrical and Computer Engineering department. At workstations, they design the integrated circuit for silicon chips. Then they mail them electronically to a silicon broker, MOSIS. At the broker's, designs from many sources are gathered on a single large silicon wafer. In that way the fabrication cost of the chips (about $400 for a tiny chip) can be shared.

Analog VLSI systems, using designs like those discussed in this article, employ the same basic devices as digital VLSI circuits. However, in analog designs, the transistors are not just used as switches. Furthermore, they operate at current levels a thousand times smaller. At these low current levels, the circuits are inherently slower and less accurate. However, these are the same constraints faced by biological information processing systems, which have evolved effective systems of cooperative behavior among many simple processing elements.

Mapping neural circuits onto silicon: The input layer of a biological retina

Figure A

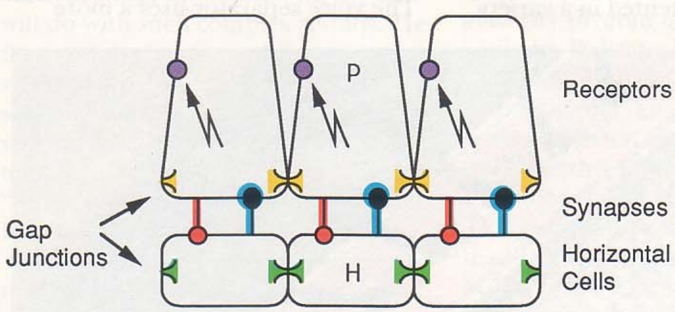

Matching colors show corresponding parts of a biological retina (simplified at top) and a silicon retina. Where it's purple, light energy enters. Red and blue stand for chemical synapses, while green and yellow synapses are electrical.

Figure B

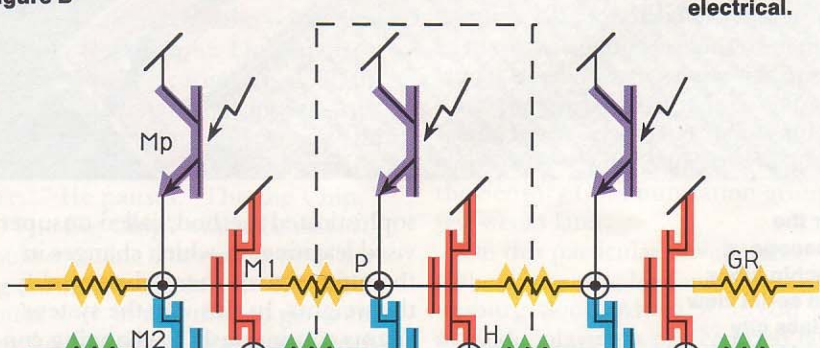

COURTESY ANDREAS ANDREOU

Figure C

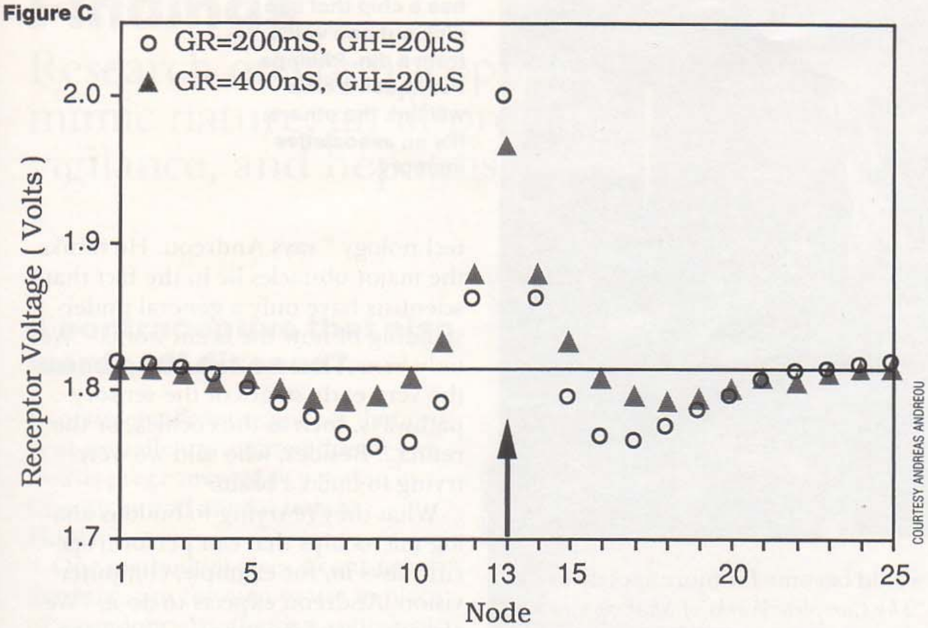

Because the chip is analog, it has many states besides simple ON (above the line) and OFF. Its "Mexican hat" response curve also occurs in biological retinas and in many parts of the brain.

uses a surprisingly simple structure to handle many functions (signal transduction, logarithmic compression, intensity normalization, and contrast enhancement), all at the same time and all in real time. This input layer is called outer-plexiform, and it consists of just three types of cells, two of which—the photoreceptors and the horizontal cells—are shown in the simplified Figure A. (The third cell type, called a bipolar cell, carries signals from the photoreceptors to the output layer of the retina.) Specialized structures, called synapses, mediate chemical and electrical interactions between the cells. The filled circles (here colored blue) are inhibitory chemical synapses, while the open ones (red) are excitatory chemical synapses. The gate-like symbols represent *electrical* synapses, which are called gap junctions; they pass electrical current between the cells. In this particular example, which does not correspond to the retina of any particular species, we have included all known interactions between the cells.

In an actual retina, cells form a densely packed two-dimensional array that lies in the back of the eye. Light enters the eye and activates photoreceptors. The photoreceptors transform light energy into current, which excites the receptors and diffuses through the gap junctions, activating neighboring receptors as well. In addition, the receptors excite the horizontal cells. They, in turn, inhibit the receptors. The key point is that excitation and inhibition take place all at the same time. Activity also spreads within the horizontal cells through gap junctions.

We often think of cells as firing on and off. However, the outer-plexiform processes information in truly analog fashion, using signals that change continuously, both over time and in voltage level.

What we have done in our laboratory is to map outer-plexiform retina processing onto silicon, as shown in Figure B. Colors show correspondences to the biological neural circuit. The silicon and biological circuits have roughly the same areas, operate at about the same power level, and have similar functionalities.

What does the silicon retina do? It gives results like the "Mexican hat" response curve in Figure C, in which signs above the line represent excitation, while signs below represent inhibition. Note how excitation of a single receptor creates a neighborhood of activation (in the middle), surrounded by inhibition. This is the well-known "center/surround organization," also found in biological retinas and in other parts of the brain.

The technical information: For chemical synapses, we use non-linear transconductances (MOS transistors in subthreshold, M1 and M2), while gap junctions are realized using active linear and non-linear conductances, GR and GH. The conductance of the silicon gap junctions is variable; it can be changed manually or can be adapted by the incident light intensity. Nodes (equipotential regions) in the top layer correspond to photoreceptors P, while those in the lower layer represent the horizontal cells H. A parasitic lateral bipolar transistor Mp is used to transduce light into current. It sources current to the receptor nodes while M2 sinks current from these nodes; these opposing effects correspond to excitation and inhibition. M1 sources current (excites) the horizontal cell nodes. The bias current Ix sets the transconductance of M1 while the photoreceptor current determines M2's transconductance. For subthreshold operation, the voltages encode photocurrents logarithmically, allowing a large dynamic range. The basic circuit (included in dotted lines) can be replicated to form a two-dimensional analog processing surface.

On the "Mexican hat" diagram, the data correspond to a 25-node one-dimensional silicon circuit. The photoreceptors supply currents of 25nA to all nodes except node 13, which receives 100nA. The two sets of data correspond to two different values of GR. As GR decreases, the activation neighborhood gets smaller. When GR is reduced to zero (absence of gap junctions between photoreceptors), the circuit exhibits a local Winner-Takes-All behavior.

For further information: D.H. Hubel, *Eye, Brain, Vision*, Scientific American Library, Book #22, 1988, and C.A. Mead, *Analog VLSI and Neural Systems*, Addison-Wesley, 1989. Also A.G. Andreou et al., "Current-Mode Subthreshold MOS circuits for VLSI Neural Systems," IEEE Transactions on Neural Networks, 1991.

BILL McALLEN

Marc Cohen (left) has a chip that can pick out one voice from a din. Philippe Pouliquen's chip will link the others: it's an associative memory.

a neural network on a single chip.

Associative memory associates two things in much the way that the *sight* of an object, say a cup sitting on a table, triggers the *name* of the object, "cup," in the human brain. Similarly, the word "cup" can be associated with the word "saucer," or the command "wash," or the warning "fragile."

Pouliquen's associative memory chip consists of 46 neurons arranged in three layers—two outside layers that handle input and output, and a central hidden layer. During supervised training the operator presents the chip with the two patterns to be associated, in this case represented by voltages, and the learning algorithms adjust the weights accordingly.

When Pouliquen presents his associative memory chip with one of seven different 16-digit numbers, it immediately spits back another of the seven numbers. Such are the humble beginnings of an engineering project. But following the next refinement, says Pouliquen, you will be able to type in "cup" and the chip will give back "saucer," "Spode," and so on.

In time, associative memories may well spawn spin-offs such as rapid-fire computer dictionaries. "A dictionary on a chip!" says Andreou, ebullient. Being analog, these chips would find definitions by simultaneously comparing all words with the original, instead of plodding through an alphabetized computer dictionary.

Dictionaries are only the beginning because, importantly, the words will not need to be in alphabetical order. Concordances and other data bases would become far more speedy to use. "*The Complete Works of Shakespeare* is not sorted," says Pouliquen, "but with associative memory you could present a word, and it would give you back all occurrences of that word."

The final aim, however, is to produce a giant associative memory with many neurons, then connect it to an artificial retina and an artificial ear. If the retina, for instance, scanned a cup, at the same time that someone said "Wash it," the memory would trigger a computer command to start the wash cycle. Associative memories are based on the neural network paradigm, and as such do not require a complete and total match. A cup that is blue, or chipped, or painted with roses—all could still trigger the wash cycle, as would a groggy, half-awake voice.

HOW LONG WILL IT BE before the team of engineers attains its goal of a perceptive computer, obedient to our slightest utterance? "At least 10 years," says Andreou. Before the team has anything resembling a truly perceptive computer, its members must refine the prototype voice separator, cochlea, retina, and associative memory. They must also come up with a method of interconnecting the parts—not just one to the next, but one fanning out to several others. And they need to improve ways of storing analog information. None of these obstacles is trivial.

In addition, problems will inevitably arise that cannot be foreseen. "But those are not going to be related to technology," says Andreou. He thinks the major obstacles lie in the fact that scientists have only a general understanding of how the brain works. "We only *know*," he says, "what happens at the very early stages of the sensory pathways, such as the cochlea, or the retina... Besides, who said we were trying to build a brain?"

What they're trying to build is analog microchips that can perform specific tasks in, for example, computer vision. Andreou expects to do it. "We already have," he crows. "These are just babies, but they work."

With future generations of such chips, the engineers predict, the computer revolution will really begin. "I see the whole gamut," says Jenkins. "Unmanned vehicles for investigating the inside of a nuclear reactor, or a sewer pipe, or a gas line. The artificial cochlea could be the beginning of a new kind of hearing aid. Or someday the silicon retina, if there was a way to tie it into the optic nerve, could help the blind."

Pouliquen joins in: "Little creatures that walk around. Pure analog computation. These things will be trained to associate lack of energy with power outlets, with some sort of need. When they start running down, they'd see a power outlet and head toward it."

Andreou picks up on the wish list: "Voice controlled telephones. Cars that find their own way. Machines that transcribe audiotapes.... Things we can't even imagine because we are constrained to think in certain ways by how computers are helping us now. When this new technology becomes available, our imaginations will just go crazy."

This work is being funded by a Research Initiation Award from the National Science Foundation and by grants from the Applied Physics Laboratory's IR/AD program. All student projects were fabricated through Educational Grants from the National Science Foundation.

Rachel Nowak writes for BioWorld *in Washington, D.C.*