

COMPUTER RECREATIONS

*A computer microscope zooms in for a look
at the most complex object in mathematics*

by A. K. Dewdney

The Mandelbrot set broods in silent complexity at the center of a vast two-dimensional sheet of numbers called the complex plane. When a certain operation is applied repeatedly to the numbers, the ones outside the set flee to infinity. The numbers inside remain to drift or dance about. Close to the boundary minutely choreographed wanderings mark the onset of the instability. Here is an infinite regress of detail that astonishes us with its variety, its complexity and its strange beauty.

The set is named for Benoit B. Mandelbrot, a research fellow at the IBM Thomas J. Watson Research Center in Yorktown Heights, N.Y. From his work with geometric forms Mandelbrot has developed the field he calls fractal geometry, the mathematical study of forms having a fractional dimension. In particular the boundary of the Mandelbrot set is a fractal, but it is also much more.

With the aid of a relatively simple program a computer can be converted into a kind of microscope for viewing the boundary of the Mandelbrot set. In principle one can zoom in for a closer look at any part of the set at any magnification [see cover of this issue and illustrations on pages 17-19]. From a distant vantage the set resembles a squat, wart-covered figure eight lying on its side. The inside of the figure is ominously black. Surrounding it is a halo colored electric white, which gives way to deep blues and blacks in the outer reaches of the plane.

Approaching the Mandelbrot set, one finds that each wart is a tiny figure shaped much like the parent set. Zooming in for a close look at one of the tiny figures, however, opens up an entirely different pattern: a riot of organic-looking tendrils and curlicues sweeps out in whorls and rows. Magnifying a curlicue reveals yet another scene: it is made up of pairs of whorls

joined by bridges of filigree. A magnified bridge turns out to have two curlicues sprouting from its center. In the center of this center, so to speak, is a four-way bridge with four more curlicues, and in the center of these curlicues another version of the Mandelbrot set is found.

The magnified version is not quite the same Mandelbrot set. As the zoom continues, such objects seem to reappear, but a closer look always turns up differences. Things go on this way forever, infinitely various and frighteningly lovely.

Here I shall describe two computer programs, both of which explore the effects of iterated operations such as the one that leads to the Mandelbrot set. The first program generated the colored illustrations appearing in this month's column. The program can be adapted to run on personal computers that have the appropriate hardware and software for generating graphics. It will create satisfying images even if one has access only to a monochrome display. The second program is for readers who, like me, need an occasional retreat from infinite complexity to the apparent simplicity of the finite.

The word "complex" as used here has two meanings. The usual meaning is obviously appropriate for describing the Mandelbrot set, but the word has a second and more technical sense. A number is complex when it is made up of two parts, which for historical reasons are called real and imaginary. These terms no longer have any special significance: the two parts of a complex number might as well be called Humpty and Dumpty. Thus $7 + 4i$ is a complex number with real part 7 (Humpty) and imaginary part $4i$ (Dumpty). The italic i next to the 4 shows which part of the complex number is imaginary.

Every complex number can be rep-

resented by a point in the plane; the plane of complex numbers is called the complex plane. To find $7 + 4i$ in the complex plane, start at the complex number 0, or $0 + 0i$, and measure seven units east and four units north. The resulting point represents $7 + 4i$. The complex plane is an uncountable infinity of such numbers. Their real parts and their imaginary parts can be either positive or negative and either whole numbers or decimal expansions.

Adding or multiplying two complex numbers is easy. To add $3 - 2i$ and $7 + 4i$, add the parts separately; the sum is $10 + 2i$. Multiplying complex numbers is only slightly more difficult. For example, if the symbol i is treated like the x in high school algebra, the product of $3 - 2i$ and $7 + 4i$ is $21 + 12i - 14i - 8i^2$. At this stage a special property of the symbol i must be brought into play: it happens that i^2 equals -1 . Thus the product can be simplified by collecting the real and the imaginary parts: it is $29 - 2i$.

It is now possible to describe the iterative process that generates the Mandelbrot set. Begin with the algebraic expression $z^2 + c$, where z is a complex number that is allowed to vary and c is a certain fixed complex number. Set z initially to be equal to the complex number 0. The square of z is then 0 and the result of adding c to z^2 is just c . Now substitute this result for z in the expression $z^2 + c$. The new sum is $c^2 + c$. Again substitute for z . The next sum is $(c^2 + c)^2 + c$. Continue the process, always making the output of the last step the input for the next one.

Strange things happen when the iterations are carried out for particular values of c . For example, here is what happens when c is $1 + i$:


first iteration, $1 + 3i$
second iteration, $-7 + 7i$
third iteration, $1 - 97i$

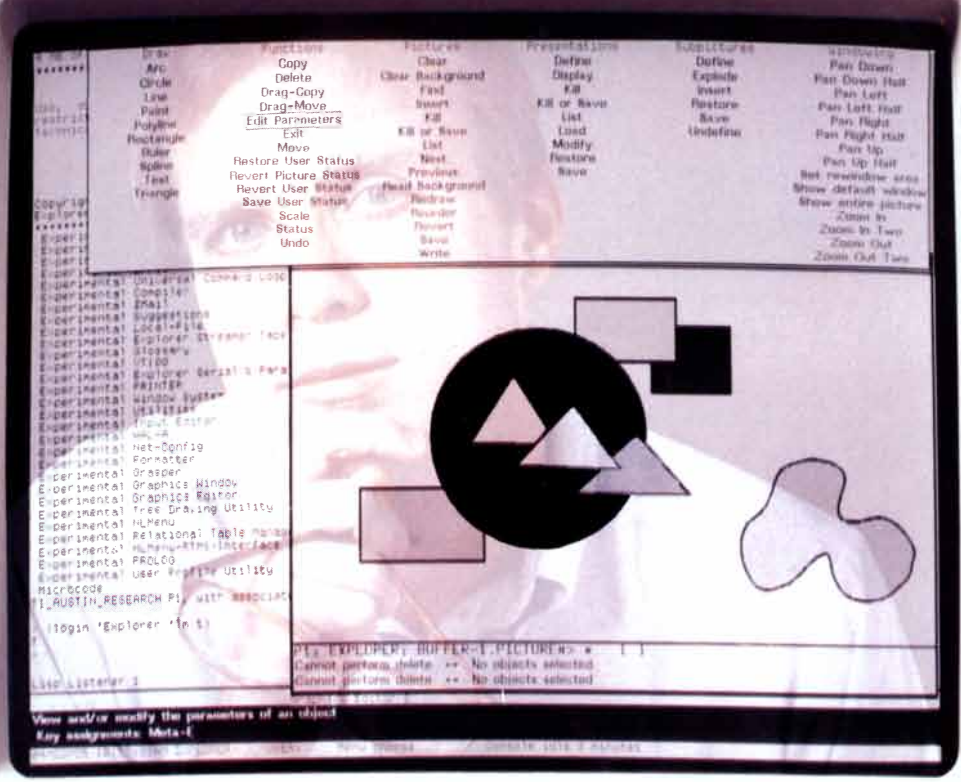
Note that the real and the imaginary parts may grow, shrink or change sign. If this process of iteration continues, the resulting complex numbers get progressively larger.

What exactly is meant by the size of a complex number? Since complex numbers correspond to points in the plane, ideas of distance apply. The size of a complex number is just its distance from the complex number 0. That distance is the hypotenuse of a right triangle whose sides are the real and the imaginary parts of the complex number. Hence to find the size of the number square each of its parts, add the two squared values and take the square root of the sum. For example, the size of the complex number

Artificial Intelligence: the vision is becoming reality— at Texas Instruments.

Whether you are already immersed in AI,
or just beginning to consider it,
this is a report of importance to you.
It touches on the highlights
of some key developments
from one of the world's most extensive,
long-term corporate R&D programs
in artificial intelligence.





TEXAS
INSTRUMENTS
EXPLORER



Bringing AI technology to personal computing.

Making technology affordable and accessible to more people is a prime mission at Texas Instruments. In putting AI research to practical use, TI has developed a number of new capabilities for personal computers that are making artificial intelligence applications available to everyone who needs them.

The Personal Consultant™—expert system development tools for personal computers.

The Personal Consultant software package from TI lets you develop and run expert systems on a TI Professional Computer or other personal computer.

Expert systems are computer programs designed to simulate the reasoning processes of human experts in a particular field—in effect, a computerized consultant. It asks a series of questions, and applies rules used by human experts to analyze the answers and make recommendations.

With Personal Consultant software, you can even ask for *help*, inquire *why* you were asked a particular question by the computer, or *how* the system reached its conclusion.

The economies offered by Personal Consultant make adaptation of this advanced technology practical for business and industry today, allowing the expansion of expert systems development and applications to a much greater degree than was affordable up to this time.

Arborist™—powerful decision-analysis software.

Arborist software illustrates the application of AI techniques to solve conventional problems. A combination of LISP processing language, multiple window displays, graphic representation of data, and plain English commands breaks through old barriers to make decision-tree problem solving techniques available to anyone who needs to make better and faster decisions.

Available as a direct result of TI research in artificial intelligence, this new software makes a versatile professional productivity tool easy-to-use and affordable on economical personal computers.

◀ *The Explorer computer system for high-performance symbolic processing.*

Natural language software—helping PCs use plain English.

A leader in the development of natural language interfaces, Texas Instruments developed NaturalLink™ software to allow people to communicate with computers in everyday English sentences, rather than in more traditional computer commands.

Invented in TI's Artificial Intelligence Research Laboratory, NaturalLink packages are available now to provide simplified access to many of today's most popular commercial PC software packages. These include word processing, spreadsheet, database, and graphics programs—programs that impact the majority of personal and business computing needs.

System developers may also build natural language interfaces for their own applications by using the TI NaturalLink Technology Package.

Computers that speak—and respond to the human voice.

Back in 1978, TI brought a major breakthrough in synthesized speech to the fingertips of children in practical, affordable, *talking* learning aids that spawned a whole new generation of electronic products.

With the emergence of personal computers, Texas Instruments has applied advanced speech technologies to PCs with the TI-Speech™ system. This innovation brings together our latest developments in both voice recognition and voice synthesis. The TI-Speech system permits users to verbally enter commands in the computer and provides sophisticated telephone management functions. It makes using a computer easier. And for some hands-busy applications, or for some handicapped individuals, it makes computerization possible.



The Business-Pro™ computer combines advanced personal computing power with capabilities for artificial intelligence applications.

At TI we're using AI today ...for tomorrow.

TI has made an extensive commitment to the application of AI technology for in-house productivity and to solve customer problems. We began our R&D efforts in 1978 and are leading the way in bringing practical, cost-effective AI products to market.

As an established world leader in the research and development of symbolic processing computers, expert systems, and natural language and speech processing, TI is now utilizing the benefits of these efforts. We've discovered new ways to apply AI solutions to improve our own productivity—to speed and improve the development process for advanced VLSI semiconductors, to increase the productivity of automated production systems, to enhance the effectiveness of vital defense systems, to analyze seismic data more accurately for more effective oil exploration, and to develop software that provides better, more timely management information.

The results of these efforts, and our own experience, is bringing thoroughly tested AI products to the marketplace

are being used to perform tasks once thought to be solely the province of human experts: planning space missions, diagnosing machinery failures, designing control systems for chemical processing plants.

Texas Instruments dedication to research in artificial intelligence has made significant contributions to this exciting new field. As one of the world's largest technology companies, TI has the resources and commitment necessary to continue its leadership role in the development and production of reliable new products which capitalize on the results of artificial intelligence research.

The next generation of symbolic processing has already begun at TI.

One such program currently underway at Texas Instruments is the development of a LISP processor on a single chip. As part of the Strategic Computing Program of the Defense

know you will stay at the leading edge of this fast developing technology—a potentially critical factor in competitive strategies for tomorrow.

TI's worldwide support and quality products help you make the most of AI technology.


To make artificial intelligence a practical reality takes more than product innovation. Texas Instruments already has a worldwide support system to provide technical assistance and follow-through. It's a total program of training, documentation, engineering support, and field service designed to provide you with the assistance needed to take maximum advantage of advanced AI technologies.

And for value, TI's exacting design and manufacturing standards enable us to deliver a higher quality product, a more reliable product, with greater functionality and at lower cost than would otherwise be possible.

Explore the promise of AI with TI.

- A proven track record for innovation in new technologies.
- An extensive, long-term R&D effort in artificial intelligence and the resources and commitment to sustain it.
- A leadership role in AI research applied to practical, affordable products—commercially available now—with new products under development.
- Consider TI's one-of-a-kind credentials in artificial intelligence.

Whether you are just beginning to think about this new wave of computing technology or already have a program in place, Texas Instruments can help you lead the way.



A new single-chip LISP processor, being developed by Texas Instruments, will have 10 times the power of previously available symbolic processors.

and making them available to an ever increasing number of disciplines.

Expert systems, natural language processing, and speech recognition have been applied to fields ranging from genetic engineering to financial management to education. Symbolic processing computers such as Explorer

Advanced Research Projects Agency, this microprocessor is designed to provide up to 10 times the power of symbolic processors previously available.

This new chip will lead the way to many new commercial applications of artificial intelligence. When you link your AI development work to TI, you

For more information.

If you would like to know more about the exciting new developments in AI at TI, write to us at Texas Instruments, P.O. Box 809063, Dept. DEE02, Dallas, TX 75380-9063 and we'll send you additional information.

The next generation of computer technology is coming of age— at Texas Instruments.

Once again, Texas Instruments is pioneering important new technology which promises dramatic increases in human productivity. And, once again, TI brings technology within your reach through practical, problem-solving products at affordable prices.

Artificial intelligence is an emerging set of technologies whose goal is to enable computers to solve problems traditionally thought to require human intelligence or capabilities.

Key to these next generation systems is a significant new approach to computer programming known as symbolic processing. Conventional computer programs require precise sequences of mathematical steps carried out in a prestructured manner. Symbolic processing works with ideas and knowledge rather than numbers, analogous to the way humans reason with knowledge they possess.

Conventional computers experience many shortcomings in areas where symbolic processing systems have proven especially effective: dealing with complex problems, interpreting information, using "rules of thumb" gained by experience, and handling uncertain or incomplete information.

SYMBOLIC vs. ARITHMETIC

Qualitative	Quantitative
Logical	Numerical
Inferential	Computational

People conceptualize ideas in symbolic terms rather than numbers. Symbolic processing computers manipulate information and knowledge in much the same way.

Symbolic processing, combined with other AI technologies—natural language, speech recognition and synthesis, computer vision—promises to open new dimensions in the way computers serve people. Using AI, computers can now be applied to the broader range of problem solving and decision making that people continually face in the real world.

At TI we're making AI work for you—now.

The transition from basic research to



The Explorer system represents an important advancement in providing exploratory programming and rapid prototyping capabilities for faster development of AI applications.

useful products is a commitment at Texas Instruments. We're concentrating not only on fundamental AI research, but on developing and marketing useful products that take full advantage of our experience. Today, TI has a number of products that are a direct result of this commitment.

The TI Explorer™ system. Solving problems beyond the reach of traditional computing.

The Explorer system from Texas Instruments uses a new computer architecture especially designed for AI applications. It executes the language of artificial intelligence, LISP (List Processing). Programmer-friendly ease-of-use features promote rapid productivity for new users.

In contrast to conventional computers, which are designed for numerical data processing, the Explorer system is

designed for the efficient processing of symbols and concepts, which represent real-world objects, their properties and relationships, and "chunks" of knowledge.

Until recently, symbolic processing systems were large, expensive computers that required special environments. This limited their use primarily to research laboratories. But, using its advanced semiconductor capability, TI has set a new standard for symbolic processing computers. The Explorer system's compact design, advanced user console, and powerful software packages now make it possible to move artificial intelligence solutions out of the laboratory and into the workplace.

For the company working on AI development, or for the organization just beginning to evaluate it, The Explorer system from Texas Instruments is a "must" consideration.

Discover new dimensions in computing with the Explorer symbolic processing system.

From the company that pioneered the transistor radio, integrated circuit, electronic calculator, integrated-circuit computer, computer-on-a-chip, micro-processor, and synthetic speech.

Once again, Texas Instruments is pioneering an important new technology and bringing it within your reach through a practical, problem-solving product at an affordable price. The Explorer computer system will play a vital role both in extending the boundaries of knowledge and in delivering the products emerging from artificial intelligence research.

A closer look at the Explorer system.

The Explorer system is an advanced computer system designed for high-performance symbolic processing. The Explorer is ideally suited for the development and execution of software that employs artificial intelligence techniques to help solve complex application problems.

The LISP environment.

The system features one of the most productive software development environments available today. The software is based on Common LISP to promote portability and consistency of software applications among different LISP machines. The Explorer also provides high-level extensions to LISP, including Flavors—an object-oriented programming facility.

Unique software.

In addition to the standard LISP machine environment, Texas Instruments provides a number of unique software packages as a standard part of the system software. The Command Interface Toolkit provides standard interfaces to the system. The Suggestions Menu System helps novices



rapidly learn the Explorer environment. The Glossary Utility offers online definitions of terms. These packages greatly reduce the time it takes a new user to become productive on the Explorer system.

Advanced hardware.

The Explorer hardware supports high-speed symbolic processing through a number of advanced architectural features. These include a tagged

architecture for run-time data typing, bit-field hardware for manipulating complex data structures, hardware assisted memory management (garbage collection), and a 128M byte virtual address space. In addition, physical memory can be expanded to 16M bytes and disk storage can expand to 1120M bytes unformatted (896M bytes formatted).

Explorer, Personal Consultant, Arborist, NaturalLink, TI-Speech, and Business-Pro are trademarks of Texas Instruments Incorporated.

TEXAS 
INSTRUMENTS

Creating useful products
and services for you.

$7 + 4i$ is the square root of $7^2 + 4^2$, or approximately 8.062. When complex numbers reach a certain size under the iterative process I have just described, they grow very quickly: indeed, after a few more iterations they exceed the capacity of any computer.

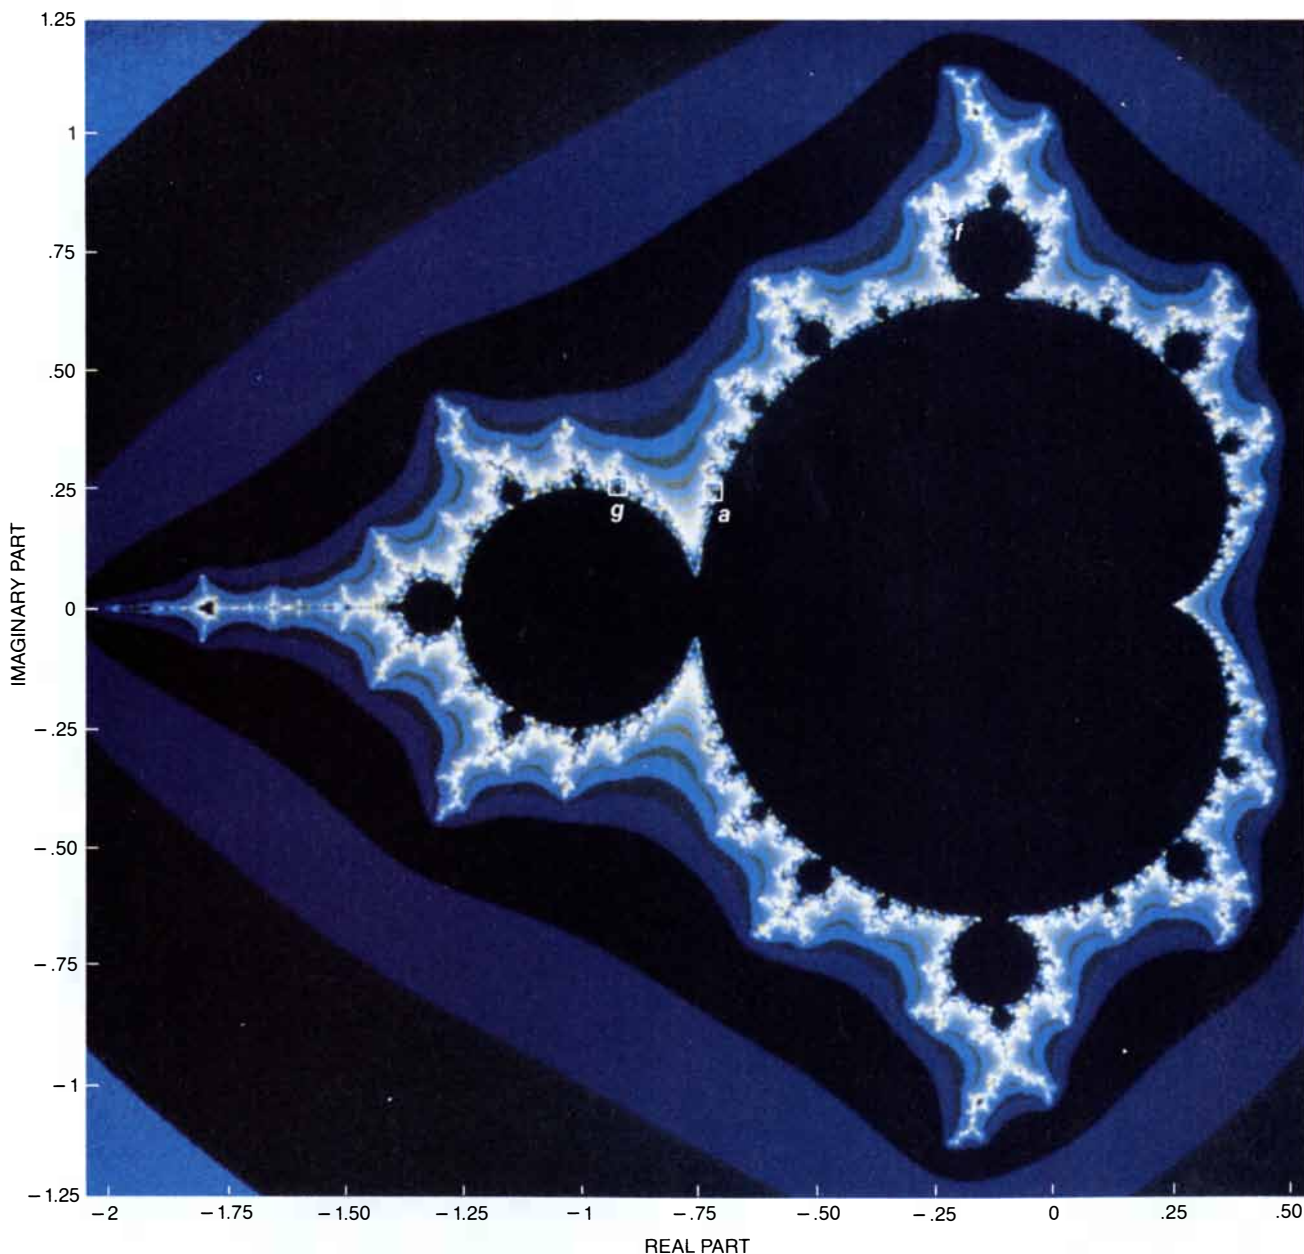
Fortunately I can ignore all the complex numbers c that run screaming off to infinity. The Mandelbrot set is the set of all complex numbers c for which the size of $z^2 + c$ is finite even after an indefinitely large number of iterations. The program I am about to describe searches for such numbers. I am indebted in all of this to John H. Hubbard, a mathematician at Cornell Uni-

versity. Hubbard is an authority on the Mandelbrot set, and he was one of the first people to make computer-generated images of it. Most of the images in this article were made by Heinz-Otto Peitgen and his colleagues at the University of Bremen. Peitgen learned the art from Hubbard.

Hubbard's program has inspired a program I call MANDELZOOM. The program sets up an array called *pic*, which is needed for saving pictures. The entries of *pic* are separate picture elements called pixels, which are arranged in a grid pattern. Hubbard's array has 400 columns and 400 rows,

and Peitgen's is even larger. Readers who want to adapt MANDELZOOM for personal use must choose an array suited to their equipment and temperament. Larger arrays impose a longer wait for the pictures, but they improve the resolution.

In the first part of MANDELZOOM one may select any square region of the complex plane to be examined. Specify the southwest corner of the square with the complex number to which it corresponds. Two variables in the program, *acorn* and *bcorn*, enable one to enter the real part and the imaginary part of the number respectively. Specify the length of each side of the



The Mandelbrot set and its coordinates in the complex plane. The details shown on the cover and on the next two pages are outlined

square by entering a value for a variable called *side*.

The second part of the program adjusts the array *pic* to match the square of interest by computing the size of a variable called *gap*. *Gap* is the distance within the square, between adjacent pixels. To obtain *gap* divide *side* by the number of rows (or columns) in *pic*.

The heart of the program is its third part. Here a search is made for the complex numbers *c* in the Mandelbrot set, and colors are assigned to the numbers that are, in a special sense, nearby. The procedure must be carried out once for every pixel; thus Hubbard's 400-by-400 array requires 160,000 separate computations. Assume the program is currently working on the pixel in row *m* and column *n*; the third part then breaks down into four steps:

1. Calculate one complex number *c*

that is assumed to represent the pixel: add $n \times \text{gap}$ to *acorn* to obtain the real part *ac* of *c*; add $m \times \text{gap}$ to *bcorn* to obtain the imaginary part *bc* of *c*. It is not necessary to include the imaginary number *i* in the program.

2. Set a complex variable *z* (which has parts *az* and *bz*) equal to $0 + 0i$. Set an integer variable called *count* equal to 0.

3. Carry out the following three steps repeatedly, until either the size of *z* exceeds 2 or the size of *count* exceeds 1,000, whichever comes first:

$$\begin{aligned} z &\leftarrow z^2 + c \\ \text{count} &\leftarrow \text{count} + 1 \\ \text{size} &\leftarrow \text{size of } z \end{aligned}$$

Why is the number 2 so important? A straightforward result in the theory of complex-number iterations guarantees

that the iterations will drive *z* to infinity if and only if at some stage *z* reaches a size of 2 or greater. It turns out that relatively many points with an infinite destiny reach 2 after only a few iterations. Their slower cousins become increasingly rare at higher values of the variable *count*.

4. Assign a color to *pic* (*m,n*) according to the value reached by *count* at the end of step 3. Display the color of the corresponding pixel on the screen. Note that the color of a pixel depends on only one complex number within its tiny domain, namely the one at its northeast corner; the behavior of this number then represents the behavior of the entire pixel.

The scheme for assigning colors requires that the range of *count* values attained within the array be grouped into subranges, one subrange for each



Successive enlargements of the "shepherd's crook" in region a of the image on the preceding page


color. Pixels for which the size of z reaches 2 after only a few iterations are colored red. Pixels for which the size of z reaches 2 after relatively many iterations are colored violet, at the other end of the spectrum. Pixels for which the size of z is less than 2 even after 1,000 iterations are assumed to lie in the Mandelbrot set; they are colored black.

It makes sense to leave the colors unspecified until the range of *count* values in a particular square has been determined. If the range is narrow, the entire color spectrum can then be assigned within that range. Thus Hubbard suggests that in step 4 only the value of *count* be assigned to each array element of *pic*. A separate program can then scan the array, determine the high and low values of *count* and assign the spectrum accordingly. Readers who get this far will certainly find workable schemes.


The reader who does not have a color monitor can still take part in black and white. Complex numbers for which z is larger than 2 after r iterations are colored white. The rest are colored black. Adjust r to taste. To avoid all-night runs the array can be, say, 100 rows by 100 columns. Hubbard also suggests it is perfectly reasonable to reduce the maximum number of iterations per point from 1,000 to 100. The output of such a program is a suggestive, pointillistic image of its colored counterpart [see illustration on next page].

How powerful is the "zoom lens" of a personal computer? It depends to some degree on the effective size of the numbers the machine can manipulate. For example, according to Magi (my microcomputer amanuensis at the University of Western Ontario), the IBM PC uses the 8088 microprocessor, a chip manufactured by the Intel Corporation designed to manipulate 16-bit numbers. A facility called double precision makes it possible to increase the length of each number to 32 bits. With such double precision Magi and I calculate that magnifications on the order of 30,000 times can be realized. Higher precision software that in effect strings these numbers together can enhance the numerical precision to hundreds of significant digits. The magnification of the Mandelbrot set theoretically attainable with such precision is far greater than the magnification needed to resolve the nucleus of the atom.


Where should one explore the complex plane? Near the Mandelbrot set, of course, but where precisely? Hubbard says that "there are zillions of beautiful spots." Like a tourist in a



A compound eye peering out from a region of image d on the opposite page



A miniature Mandelbrot in region f on page 17, tethered to the main set by a filament



Pointillist, miniature Mandelbrot generated by a monochrome monitor

land of infinite beauty, he bubbles with suggestions about places readers may want to explore. They do not have names like Hawaii or Hong Kong: "Try the area with the real part between .26 and .27 and the imaginary part between 0 and .01." He has also suggested two other places:

<i>Real Part</i>	<i>Imaginary Part</i>
-.76 to -.74	.01 to .03
-1.26 to -1.24	.01 to .03

The reader who examines the color images accompanying this article should bear in mind that any point having a color other than black does not belong to the Mandelbrot set. Much of the beauty resides in the halo of colors assigned to the fleeing points. Indeed, if one were to view the set in isolation, its image might not be so pleasing: the set is covered all over with filaments and with miniature versions of itself.

In fact none of the miniature Mandelbrots are exact copies of the parent set and none of them are exactly alike. Near the parent set there are even more miniature Mandelbrots, apparently suspended freely in the complex plane. The appearance is deceiving. An amazing theorem proved by Hubbard and a colleague, Adrian

Douady of the University of Paris, states that the Mandelbrot set is connected. Hence even the miniature Mandelbrots that seem to be suspended in the plane are attached by filaments to the parent set. The minatures are found almost everywhere near the parent set and they come in all sizes. Every square in the region includes an infinite number of them, of which at most only a few are visible at any given magnification. According to Hubbard, the Mandelbrot set is "the most complicated object in mathematics."

Readers with a simple appetite for more color images of the Mandelbrot set and other mathematical objects can write to Hubbard for a brochure (Department of Mathematics, Cornell University, Ithaca, N.Y. 14853). The brochure includes an order form with which one can buy 16-inch-square color prints that are similar in quality to the Peitgen images shown here.

Confronted with infinite complexity it is comforting to take refuge in the finite. Iterating a squaring process on a finite set of ordinary integers also gives rise to interesting structures. The structures are not geometric but combinatorial.

Pick any number at random from 0 through 99. Square it and extract the

last two digits of the result, which must also be a number from 0 through 99. For example, 59^2 is equal to 3,481; the last two digits are 81. Repeat the process and sooner or later you will generate a number you have already encountered. For example, 81 leads to the sequence 61, 21, 41 and 81, and this sequence of four numbers is then repeated indefinitely. It turns out that such loops always arise from iterative processes on finite sets. Indeed, it is easy to see there must be at least one repeated number after 100 operations in a set of 100 numbers; the first repeated number then leads to a loop. There is a beautiful program for detecting the loops that requires almost no memory, but more of this later.

It takes only an hour to diagram the results of the squaring process. Represent each number from 0 through 99 by a separate point on a sheet of paper. If the squaring process leads from one number to a new number, join the corresponding points with an arrow. For example, an arrow should run from point 59 to point 81. The first few connections in the diagram may lead to tangled loops, and so it is a good idea to redraw them from time to time in such a way that no two arrows cross. A nonintersecting iteration diagram is always possible.

One can go even further. Separate subdiagrams often arise, and they can be displayed in a way that highlights some of the symmetries arising from the iterations. For example, the nonintersecting iteration diagram for the squaring process on the integers from 0 through 99 includes six unconnected subdiagrams. The pieces come in identical pairs and each piece is highly symmetrical [see illustration on opposite page]. Can the reader explain the symmetry? What would happen if the integers from 0 through 119 were used instead? Is there a relation between the number of unconnected pieces found in the diagram and the largest integer in the sequence?

Similar patterns of iteration hold for some of the complex numbers in the Mandelbrot set: for certain values of c repeated iterations of $z^2 + c$ can lead to a finite cycle of complex numbers. For example, the complex number $0 + 1i$ leads to an indefinite oscillation between the two complex numbers $-1 + 1i$ and $0 - 1i$. The cycle may even have only one member. Whether such cycles are found in a finite set or in the infinite Mandelbrot set, they are called attractors.

Each of the six parts of the iteration diagram for the integers 0 through 99 includes one attractor. Geometrically the attractor can be represented as a polygon, and the sets of numbers that



THE SKY WAS THE LIMIT.

AT&T has shattered the information barrier — with a beam of light.

Recently, AT&T Bell Laboratories set the world record for transmission capacity of a lightwave communications system — 20 billion pulses of light per second. The equivalent of 300,000 conversations, sent 42 miles, on a hair-thin fiber of super-transparent glass. But that's really getting ahead of the story.

Actually, the 20-gigabit record is only one of a series of AT&T achievements in the technology of lightwave communications.

But what does that record mean?

The Light Solution To A Heavy Problem

All of us face a major problem in this Information Age: too much data and too little information. The 20-gigabit lightwave record means AT&T is helping to solve the problem.

For data to become useful information, it must first be quickly, accurately and securely moved to a data transformer — a computer, for instance. Getting there, however, hasn't always been half the fun.

Metallic pathways have a limited transmission speed, sensitivity to electrical interference and potential for interception — factors that reduce the effectiveness of today's powerful computers. Factors that are eliminated by lightwave communications technology.

Ten Goes Into One 20 Billion Times

Three primary components make up any lightwave communications system. On the transmitting end, a laser or light-emitting diode; on the receiving end, a highly sensitive photodetector; and in the middle, super-transparent glass fibers we call lightguides.

Installing these fibers is a major cost of a lightwave communications

system. So, once installed they should stay put — increased capacity should come from fibers carrying more, rather than from more fibers.

Which brings us to the 20-billion bit-per-second story — about experimental technology that has the potential to upgrade installed fiber to meet any foreseeable capacity needs.


Using new, sophisticated lightwave system components, we multiplexed (combined) the outputs from 10 slightly different colored 2-billion bit-per-second laser beams into a single 20-billion bit-per-second data stream.

Playing Both Ends Against The Middle

But, let's start at the beginning — the 10 distributed feedback laser transmitters.

These powerful semiconductor lasers can be grown to produce light of different, but very precise, wavelengths. The lasers we used transmitted in the 1.55 micron (infrared) range, with only minuscule fractions of a micron between their wavelengths. The purity and stability of the beams let us pack their ten colors into the most efficient transmitting region of our single-mode, silica-core fiber.

To make the original 10 beams into one, a fiber from each laser was fed into a new lightwave multiplexer — a



20-gigabit
multiplexer

prism-like grating that exactly aimed each beam into the single transmission fiber. Over 42 miles later, a second grating fanned the beam back into its original 10 colors for delivery to 10 exceptionally sensitive avalanche photodetectors — receivers that convert the light pulses back into electrical signals and amplify them many times.

A similar avalanche photodetector

was the receiver when AT&T Bell Laboratories set the world record for unboosted lightwave transmission — 125 miles at 420 million bits per second.

From Sea To Shining Sea

System capacity is important. But system reliability is vital. Especially when the system is going under 10 thousand miles of water — and is expected to last for 25 years.

AT&T is going to build the first lightwave communications system under the Atlantic Ocean. A similar system is planned for the Pacific. In 1988, laser beams traveling through two pairs of glass fibers will carry the equivalent of 37,800 simultaneous conversations overseas, underwater, from the U.S. to Europe and the Far East.

AT&T has manufactured and installed lightwave systems — as large as the 780-mile Northeast Corridor and as small as single-office local area networks — containing enough fiber to stretch to the moon and back. And the capacity of each network is tailored to meet the unique needs of its users.

Systems being installed in 1985 will be able to grow from 6,000 up to 24,000 simultaneous conversations on a single pair of fibers.

AT&T is meeting today's needs with lightwave systems that are growable, flexible and ultra-reliable. And anticipating tomorrow's needs with a whole spectrum of leading-edge lightwave communications technologies.



AT&T

The right choice.

readily adapted to computer programs.

The program is called RHOP because the sequence of numbers that eventually repeats itself resembles a piece of rope with a loop at one end. It also resembles the Greek letter rho (ρ). There are two variables in the program called *slow* and *fast*. Initially both variables are assigned the value of the starting number. The iterative cycle of the program includes just three instructions:

$$\begin{aligned} fast &\leftarrow fast \times fast \pmod{100} \\ fast &\leftarrow fast \times fast \pmod{100} \\ slow &\leftarrow slow \times slow \pmod{100} \end{aligned}$$

The operation mod 100 extracts the last two digits of the products. Note that the squaring is done twice on the number *fast* but only once on the number *slow*. *Fast* makes its way from the tail to the head of the rho twice as fast as *slow* does. Within the head *fast* catches up with *slow* by the time *slow* has gone partway around. The program exits from its iterative cycle when *fast* is equal to *slow*.

The attractor is identified by reiterating the squaring process for the number currently assigned to *slow*. When that number recurs, halt the program and print the intervening sequence of numbers.

I should be delighted to see readers' diagrams that explore the effects of iterative squaring on finite realms of varying size. The diagrams can be done on a computer or by hand. Discrete iteration is a newly developing mathematical field with applications in computer science, biomathematics, physics and sociology. Theorists might watch for a book on the subject by François Robert of the University of Grenoble.

The two-dimensional beings who inhabit the planet Arde are deeply grateful to the many readers who tried to improve the crossover circuit I described in May. That circuit is made up of 12 two-input *nand*-gates. I asked readers to find the minimum number of *nand*-gates—and *nand*-gates only—from which a crossover circuit can be built. Most of the circuits submitted have 10 gates, a mild improvement, but three readers found an eight-gate crossover [see illustration on this page].

In the eight-gate circuit there is one three-input *nand*-gate and two single-input *nand*-gates. The latter act as inverters, converting a 0 signal into a 1 signal and vice versa. The three readers who discovered the eight-gate solution are Eric D. Carlson of Cambridge, Mass., Dale C. Koepp of San Jose, Calif., and Steve Sullivan of Beaverton, Ore. I have passed their names

along with the improved crossover circuit to my Ardean friends. Believe it or not, the same crossover circuit appears under U.S. Patent 3,248,573 (April 26, 1966). Robert L. Frank, who is a systems consultant in Birmingham, Mich., wrote that the patent was awarded to Lester M. Spandorfer of Cheltenham, Pa., Albert B. Tonik of Dresher, Pa., and Shimon Even of Cambridge, Mass.


It seems natural to wonder whether the circuit actually appears in any present-day device. It is also natural to wonder whether there is an even smaller *nand* crossover. One supposes not.

C. Walter Johnson of Long Beach, Calif., wrote to me describing a wide variety of planar circuits that incorporate several types of gate. Apparently it is possible to build not only crossover circuits in two dimensions but also planar flip-flops. The flip-flops provide memory for a two-dimensional computer.

One-dimensional computers in the form of cellular automata have been investigated by Stephen Wolfram of the Institute for Advanced Study in Princeton. It is too early to say what contributions readers may have made to this field after reading Wolfram's "Glider Gun Guidelines," but I can pass along some initial reactions. A sin of commission sent a few readers off chasing gliders in the line automaton code-numbered 792. Wolfram and I meant to specify code 357. A sin of omission was my decision not to mention the line automata known to be capable of universal computation. I thought of describing such a line automaton, first constructed by Alvy Ray Smith in 1970. At the time Smith was a graduate student at Stanford University. I was afraid that the description of Smith's universal line automaton would unduly complicate the article: the automaton has 18 states ($k = 18$) and three-cell neighborhoods ($r = 1$).

Arthur L. Rubin of Los Angeles has made a sensible suggestion for defining the speed of light in an arbitrary line automaton. Rubin's suggestion corrects a defect in an earlier definition that sets the speed of light equal to one cell per unit of time. The old definition ignores the possibility that not all automata can attain such speeds. The revised speed of light is "the maximum speed of propagation of any impulse (say to the right)." The leading edge of the impulse is defined by the condition that only 0's can lie to its right. Rubin goes on to prove that the speed of light is 1/3 for the line automaton code-numbered 792.

In my May column I also asked whether the line automaton called



A crossover circuit with eight *nand*-gates

Ripple has a one-way glider gun. Gliders fired from such a gun would spew out unendingly to the right but never to the left. William B. Lipp of Milford, Conn., has made a simple and charming argument against the existence of such a gun. "Consider a pattern," he writes, "that never has nonzero values to the left of some block labeled 0. Observe that the leftmost nonzero value in the pattern must always be a 1. If it were a 2, the 2 would ripple to the left forever, thus contradicting the assumption that no nonzero entries lie to the left of block 0. But the leftmost 1 must become 0 on the next cycle, moving the left boundary of the pattern at least one block to the right." Thus either a glider ripples to the left or its gun is eaten away by 0's.

Other readers sought to show that Ripple is not capable of universal computation. For some automata one can prove a sufficient condition, namely that the halting problem is decidable. Ripple halts when all its cells contain 0's, but the halting conditions for any universal computing machine it might contain could be quite different.

Several readers attempted constructions of line automata capable of universal computation, among them Frank Adams of East Hartford, Conn., Jonathan Amsterdam of Cambridge, Mass., Kiyoshi Igusa of Brandeis University and Carl Kadie of East Peoria, Ill. The constructions are all straightforward and believable, but Kadie, not content with his one-dimensional automaton, went on to suggest a zero-dimensional one. It would consist of a single cell, and it seems reasonable to call it a point automaton. Readers with a theoretical bent might enjoy pondering the universality of a point automaton. Is it possible?

Alvy Ray Smith published his proof for the existence of a computation-universal line automaton in 1971, in *Journal of the Association for Computing Machinery*. One might have thought a career with such an auspicious start would today be blossoming in some well-known academic institution. Instead it has blossomed in a quite different setting: Smith is director of computer-graphics research for Lucasfilm, Ltd., in San Rafael, Calif. In a future column I hope to report on some of the amazing cinematic effects produced at the Lucasfilm laboratory.

1,310 PROBLEM-SOLVING PAGES IN 4 VOLUMES. EVERYTHING FROM "ARISTOTLE'S WHEEL" TO ALGORITHMS. A \$111.40 VALUE FOR ONLY \$3.95. A \$107.45 SAVINGS!

Yours when you join The Library of Science. You simply agree to buy 3 more books—at handsome discounts—within the next 12 months.

It's that easy, that inexpensive to solve your problem-solving problems. Our Selection Committee assembled four volumes into a set called Practical Tools for Problem-Solving because they all share a common concern: the systematic investigation of the most efficient procedures you can use to solve specific problems. By implementing methods expounded in these books, you can solve virtually any type of abstract problem in math, language and meaning, computer science, physics, etc.

And remember, you get all 4 volumes—a \$111.40 value—for only \$3.95. A savings of \$107.45.

A. Wheels, Life, and Other Mathematical Amusements by Martin Gardner. Join the author in the cosmic fun by testing your wits against mathematical puzzles, paradoxes, and other mathematical recreations. The problems represent such diverse topics as: diophantine analysis, probability paradoxes, and combinatorial

theory. Publisher's Price: \$16.95.

B. Bypasses by Z. A. Melzak. Subtitled "A Simple Approach to Complexity," this volume offers a new and unique approach to problem-solving and demonstrates its use in a wide variety of scientific and engineering concerns. Publisher's Price: \$40.50.

C. How To Solve It, second edition, by George Polya. Any problem can be solved scientifically. But how? The answer lies in the skill of heuristics—defined as the study of how problems are solved. In HOW TO SOLVE IT, mathematical expert G. Polya spells out the methods you can use to solve any problem. Regardless of the subject at hand, Polya's methods can help you quickly arrive at a solution. Publisher's Price: \$25.00.

D. Algorithms by Robert Sedgewick. Surveys the most important algorithms and data structures in use on computers today. Covers seven major areas: mathematical algorithms, sorting, searching, string processing, geometric algorithms, graph algorithms, and advanced topics. Publisher's Price: \$28.95.

The Library of Science is the nation's oldest, most respected science book club.

Dedicated to serving both professionals and serious amateurs, it sur-

veys the entire field of published sources on scientific subjects to make the very finest works available to club members at substantial savings.

Membership Benefits:

In addition to getting Practical Tools for Problem-Solving for only \$3.95 when you join, you keep saving substantially on the books you buy. Also, you will immediately become eligible to participate in our Bonus Book Plan, with savings of up to 70% off publishers' prices. At 3-4 week intervals (16 times per year), you will receive The Library of Science News, describing the Main Selection and Alternate Selections, together with a dated reply card. If you want the Main Selection, do nothing, and it will be sent to you automatically. If you prefer another selection, or no book at all, simply indicate your choice on the card and return it by the date specified. You will have at least 10 days to decide. If, because of late mail delivery of the News, you should receive a book you do not want, we guarantee return postage.

If reply card has been removed, please write to: The Library of Science, Dept. 2BF9-00676, Riverside, N.J. 08075 to obtain membership information and an application. Scientific American 8/85

