# SUPERCOMPUTING AND THE TRANSFORMATION OF SCIENCE

William J. Kaufmann III
Larry L. Smarr

SCIENTIFIC
AMERICAN
LIBRARY

A division of HPHLP
New York

Library of Congress Cataloging-in-Publication Data

Kaufmann, William J.
Supercomputing and the transformation of science / William J. Kaufmann III, Larry L. Smarr.
p. cm.
Includes bibliographical references (p. ) and index.
ISBN 0-7167-5038-4
1. Supercomputers. 2. Science—Data processing. I. Smarr, Larry
L. II. Title.
QA76.88.K38 1993 92-32418
502′.85′411—dc20 CIP

ISSN 1040-3213



Printed in the United States of America

Scientific American Library
A division of HPHLP
New York

Distributed by W. H. Freeman and Company
41 Madison Avenue, New York, NY 10010
20 Beaumont Street, Oxford OX1 2NQ, England

2 3 4 5 6 7 8 9 0 KP 9 9 8 7 6 5 4 3

This book is number 43 in a series.

Supercomputing and the Transformation of Science

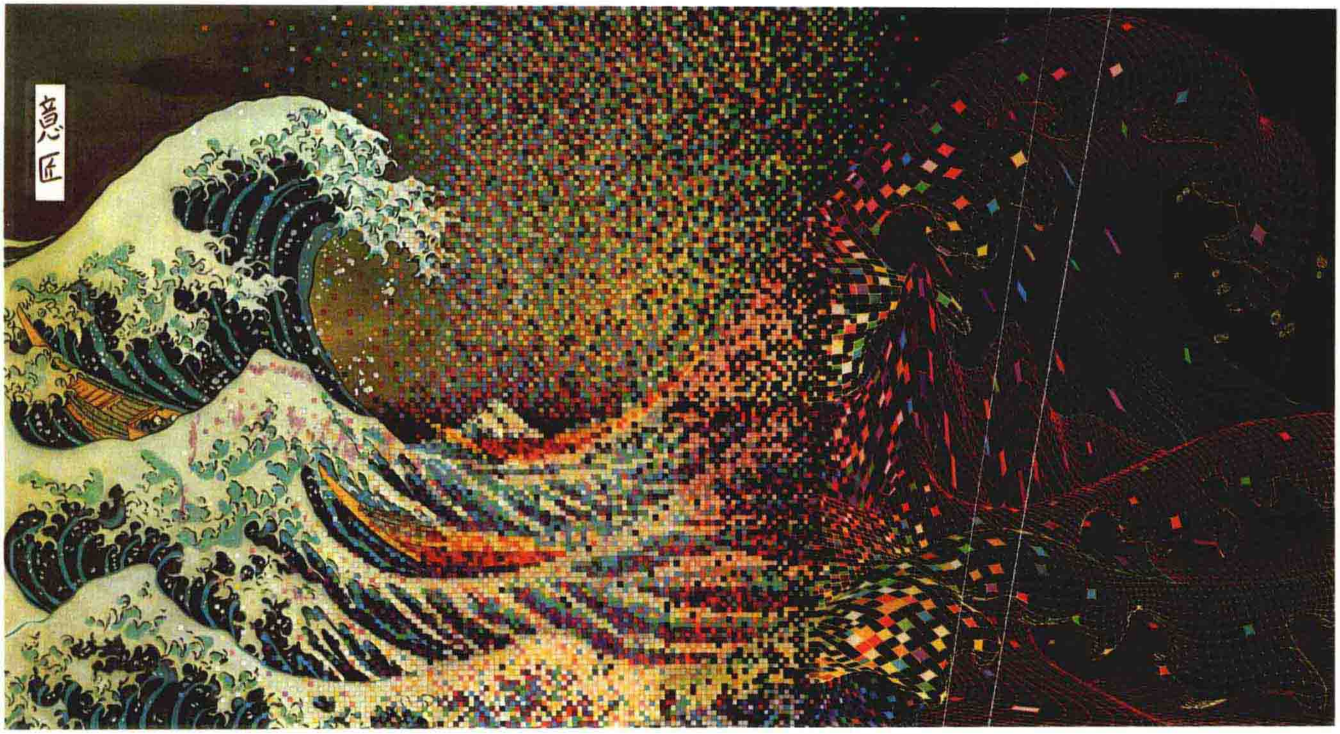


*This picture, adapted from Katsushika Hokusai's masterpiece "The Great Wave off Kanagawa," artistically displays the spirit of supercomputing. Complex phenomena, such as waves on the surface of a fluid, are modeled by covering space with a grid and then solving the laws of physics at discrete points on that grid. The finer the grid, the closer the numerical simulation is to the actual solutions of the mathematical laws of nature that govern the physical world.*

*To the intellectual giants who made*
*the digital world possible:*
*Charles Babbage, who created the computer;*
*John von Neumann, who taught us how to use it;*
*and Vannevar Bush, who foresaw how computers and*
*communications would transform science and society.*

# Preface

A new information reality, parallel to but distinct from our well-known physical reality, is emerging. In this digital reality, bits take the place of the fundamental atoms of the physical world. Scientific instruments, from telescopes to microscopes to space probes, are extending our senses to enormous distances and to tiny scales, even to other worlds. Yet, instead of producing analog photographic plates that end up being stored in a vault, modern sensors record their discoveries in digital data banks accessible to researchers anywhere over computer networks. The mathematical laws of gravity, gas dynamics, and quantum mechanics—the products of generations of scientific thought—can now be solved by digital computers to create numerical representations of the physical world of orbiting planets, thunderstorms, and new drugs. These simulations offer scientists a new means of exploring nature. Finally, our analog modes of communication by voice, print, and video are gradually being replaced by digital modes. Ultimately most of human knowledge will be stored in a common digital library.

The worldwide acceptance of the personal computer has given every desktop a window into the world of digital knowledge. Fiber optic networks are creating intricate connections between millions of desktop computers and the relatively few supercomputers, the fastest computers that exist. These same networks tie this computational infrastructure into the vast archives of data maintained by industry, government, and scientific laboratories. This shadow universe of information, the "cyberspace" of science fiction novelist William Gibson, is rapidly increasing its reach. In *Supercomputing and the Transformation of Science,* we explore the supercomputers that are the central powerhouses of this information space. Our book focuses on three themes: the evolution of supercomputers, the methodologies for using them to simulate nature, and their transformation of virtually every field of science and engineering.

The supercomputers of today run almost one trillion times faster than the fastest computer of fifty years ago. For comparison consider another transformational technology—transportation. Fifty years ago the fastest mode of transportation had a speed limit of a few hundred miles per hour, whereas today's interplanetary spacecraft attain astonishing speeds of several hundred thousand miles per hour. We can now travel a thousand times faster than we could fifty years ago, yet this speedup is only one-billionth of the increase in speed achieved by supercomputers in the same time!

Because of their speed, supercomputers can perform huge numbers of computations in a brief span of time. It is that ability that enables these machines to create simulations of the natural world. Mathematics is capable of capturing the rich phenomenology of nature, and supercomputers are capable of employing this mathematics to generate the billions of numbers necessary to simulate the behavior of natural phenomena. The rise of computer graphics has allowed these vast mountains of numbers to be translated into visual imagery that is more intuitively accessible to human beings. Throughout the book, we have used visual images instead of traditional equations to capture the essence of a technical subject.

One of our most challenging tasks as authors was to convey the fundamental idea behind most of the techniques used to adapt the mathematical equations of theoretical science for use on supercomputers: that idea is to replace the continuous world of nature with a model of that world formed of discrete units. This can be accomplished by a variety of methods: we can approximate a fluid by dividing the space it flows through into a large number of small boxes, we can represent an engineering device by a finite set of subelements, or we can simulate a galaxy as a large number of gravitating particles. As we move from the more familiar world of classical physics into the strange world of quantum physics, we show how less intuitive approaches are used to make the complexity of quantum mechanical systems computationally tractable.

After clarifying the methodology of supercomputing, we take the reader on a comprehensive tour of its frontier applications. We start our voyage of discovery with the world of the quantum and end with the cosmos, following a course that parallels the organization of physical reality into a hierarchy of levels of increasing physical scale. In this hierarchy, atoms build molecules, which in turn build cells or bulk materials, which in turn build bodily organs or manufactured objects or geological structures. As our book progresses, we move up the hierarchy: thus, in between the extremes of the atom and the whole universe, we explore the worlds of biology, engineering, and the environment.

The virtual worlds simulated by supercomputers can be probed and measured much more thoroughly than the physical world. Our exploration of these worlds is already leading to deeper understandings, and it is also turning the supercomputer into an instrument of engineering design. The result is the creation of safer, cheaper, and more reliable products in less time than would otherwise be possible. Following its adoption a decade ago in the petroleum, automobile, and aerospace industries, supercomputing is now radically altering industries producing chemicals, pharmaceuticals, consumer products, and financial services. Furthermore, supercomputers are beginning to offer one of the few techniques for rationally analyzing the problems confronting our environment. Pollution, exhaustion of natural resources, ozone depletion, and global warming are but a few of the research frontiers explored by supercomputing.

As we approach the next century, we cannot help but wonder how supercomputing will transform our ability to manipulate the physical world. Simulations are giving us detailed knowledge of materials and of biological molecules. As we acquire this knowledge, we will undoubtedly be able to restructure the traditional substances and organisms found in nature today. As we learn how to digitally control the engineering process from design through manufacturing, we will radically enhance our abilities to convert basic materials into finished products. As we gain detailed and verified models of the ecological and environmental systems of the Earth, we will have vastly greater powers for altering our world. Our hope is that the ability of computers to assist human beings in understanding the complexities of our world will lead to a growth of wisdom adequate to the challenges created by these new technologies.

We would like to thank the thousands of researchers whose pioneering use of supercomputers inspired us to write this book. We have been able to include the work of only a few rep-

resentatives from each discipline, and we apologize to all those whose research we were unable to cover. The researchers whose work we do describe often sent us materials or read drafts for accuracy on very short notice. We would like to single out those who took on the extra effort of reviewing entire sections in order to minimize the misrepresentations that are unavoidable in a broad survey such as this one. These people include Fouad Ahmad, David Ceperley, Robert Chervin, Art Freeman, Bruce Hannon, Michael Heath, Michael Fainan, Karl Hess, Eric Jakobsson, Radha Nandkumar, Michael Norman, David Pines, Michael Schlesinger, Harrell Sellers, Shankar Subramaniam, Robert Sugar, Warren Washington, Robert Wilhelmson, and Carl Woese. Finally, Michael Norman and Robert Wilhelmson put in extra effort to create original images and illustrations for our book.

Larry Smarr also benefitted from the generous help of a number of his fellow directors of supercomputing facilities, including Bill Buzzbee, Sid Karin, Mal Kalos, Michael Levine, Ralph Roskies, and Vic Peterson. The documentation staffs of these centers and the science writers for *Cray Channels* and *Supercomputing Review* made our job of identifying and developing stories about individual research efforts much easier than it would have been otherwise.

Bill Kaufmann, who traveled to numerous conferences, symposia, and supercomputing facilities throughout the world to gather material for this book, would like to thank his many friends and colleagues for their help and hospitality. Bill is especially grateful to the staffs at Los Alamos National Laboratory, Lawrence Livermore National Laboratory, NASA Ames Research Center, the Minnesota Supercomputing Institute, Sandia National Laboratory, the National Center for Supercomputing Applications, the National Center for Atmospheric Research, the European Centre for Medium-Range Weather Forecasts, the San Diego Supercomputing Center, and Cray Research, Inc.

The manuscript for this book demanded elaborate preparation, made necessary by the enormous scope of the subject matter. The support of the staff of the National Center for Supercomputing Applications, particularly in Applications, Documentation, and Media Services, was critical in many ways, as was NCSA's support structure provided by the National Science Foundation, the State of Illinois, the University of Illinois at Urbana-Champaign, and NCSA's corporate sponsors. Jim Bottum, Deputy Director of NCSA, kept the center running while the Director was busy writing. The Director's assistant, Janus Wehmer, and her secretary, Linda Griffet, provided support above and beyond the call of duty; without their dedicated efforts, the publication schedule for this book could never have been met.

It has taken five years to bring this book to fruition. During that time, the staff of the Scientific American Library has been very patient and enormously competent. From the beginning, publisher Jerry Lyons provided continual encouragement. Once we had completed our manuscript, Susan Moran provided meticulous and invaluable editing. Indeed, she deserves to be a third author of this book!

The Scientific American Library production staff did a wonderful job, considering the hundreds of color images that had to be collected and processed. We would especially like to acknowledge Larry Marcus for photo research, John Hatzakis for page layout, Alice Fernandes-Brown for our book's design, Tina Hastings for her work as project editor, Christine McAuliffe for her oversight of the line illustrations, and Sheila Anderson for her coordination of typesetting and printing.

Finally, Larry Smarr would like to express his deep appreciation for the unending patience of his wife, Janet, and his boys, Joseph and Benjamin, during the final period of preparation of the manuscript.

*William J. Kaufmann III*
*Larry L. Smarr*
*September 1992*

Supercomputing and the Transformation of Science

# Contents

1

# The Emergence of a Digital Science

Computers have a significant impact on daily life. They keep track of our telephone calls, compute our income taxes, and tally our charge cards. But one kind of computer, the supercomputer, is actually changing the way in which we conduct scientific research and engineering design. These computers, the most powerful that can be built, allow us to replace the physical world with a digital reality in the form of an array of numbers. When transformed into images on a computer screen, these numbers are easily seen to replicate essential features of phenomena in the natural world.

*In this image from a supercomputer simulation, the paths of blue and orange tracer particles give a comprehensive view of air motions in and around a developing severe thunderstorm. The image maps the cloud's raindrop content in greens and yellows at a horizontal plane 2.25 kilometers above the ground. Below 2.25 kilometers the surface enclosing rain and cloud drops is rendered as a solid and above as a transparent curtain.*

For instance, when astrophysicists wondered whether an exotic arrangement of stars could be two colliding galaxies, they programmed a computer to calculate the behavior of two galaxies as their paths crossed. The computer generated numbers, which when transformed into images, showed whether two colliding galaxies would look like the astronomical phenomenon. Other numerical models can reveal the damage from a car crash, the action of a drug on a cancer cell, the growth of a thundercloud, or the evolution of the universe.

Our ability to create these numerical models rests on the amazing fact that mathematics can be used to represent the physical universe. No one really understands why this should be so, but much of the success of modern science and engineering is based on our ability to create an abstract mapping between the motions of matter and symbols on paper. Following this approach, scientists have gradually discovered a set of mathematical equations, generally referred to as the laws of nature, that describe the physical world. These fundamental equations are written in the language of calculus, a branch of mathematics that deals with rates of change—how one quantity varies with respect to other quantities such as time or location. Such equations are precise because calculus divides time and space into points that are infinitesimally close to each other.

A developing thunderstorm can be viewed as nature's way of physically solving a subset of the laws that govern gas dynamics, heat transfer, and the properties of water. This physical solution provides an evolving temperature, pressure, wind speed, and wind direction at each point in space. It also specifies the phase state of water—whether the water is in vapor, liquid, or solid form—and it specifies the size, amount, and kind of water droplets or ice crystals.

Imagine recording the numerical value of each of these variables at all points in space in the volume of the atmosphere containing the storm and for all moments in time during the storm's evolution. This infinite set of numbers would, in principle, constitute a digital solution of the laws of nature. A different thunderstorm, evolving with another set of numbers, would represent a different digital solution to the same laws of nature. Thus to every universal law there are an enormous variety of possible physical solutions, each of which has a corresponding digital solution.

Although nature is continually constructing physical solutions to its laws, humans can neither experimentally nor theoretically produce completely accurate digital solutions. Our only practical approach to creating digital solutions is to use supercomputers to create a mathematically approximate solution to these equations, called a numerical solution. This process is often called simulation.

To perform a simulation, the scientist begins by choosing an appropriate subset of the fundamental equations. These equations, which are valid at every point in space and at every moment in time, are then replaced with a closely related set of equations defined only at selected points in space and selected moments in time. These "discretized" equations are programmed into the supercomputer. Instead of tackling the impossible chore of solving the laws of nature everywhere for all time, the supercomputer evaluates required quantities only at the selected points at prescribed time intervals.

Such an approach has a number of distinct advantages. First, scientists can replay a solution over and over, whereas they are able to observe most natural phenomena only once. Second, they can study an ensemble of solutions, each describing the same phenomenon but with variables of different values. For example, they might simulate many different thunderstorms in order to extract the common defining properties from the details of the solutions. Scientists can isolate a dominant subcomponent of a phenomenon and make a more finely resolved simulation of just that feature. Finally, they can determine all the values of the physical variables as these variables change through space and time.

However, the numerical solution by itself is of little use. Today's supercomputers are capable of performing one billion arithmetic operations per second, and a typical simulation runs for hours. Even a small portion of the results comprises billions of numbers. No scientist could digest the vast columns of numbers that stream from the computer programs used in simulation were they not transformed into images.

Simple arithmetic drives home the necessity for visualization techniques that display data in the form of pictures. In one second of operation, a modern supercomputer can generate one billion numbers, which if printed out in 10 columns of 50 lines apiece on each page, would require a pile of paper over 50 stories tall! As early as 1995, when supercomputers will be one thousand times faster than today, that one second of operation would produce enough printed output to rise more than 100 miles in height. For this reason, scientific visualization has become as important as supercomputers to the computational scientist.

By inserting more spatial points and shortening the time intervals, scientists can increase the accuracy of a simulation, because the discretized computation approaches the continuous coverage of space and time characteristic of the exact ideal solution. However, improving the realism of the simulation, either by adding more points and time intervals or by adding more laws of nature to the set of equations, increases computing time. It also greatly increases the amount of computation needed to transform the enlarged numerical output into visual images.

An insatiable craving for ever faster supercomputers is a direct result of the scientist's desire to produce solutions of increasingly realistic complexity during an allotted span of computer time. For this reason, scientists have enthusiastically adopted each new generation of digital electronic computers for over fifty years. During that time, computers have increased in speed by more than a billion times! And yet today's scientists are just as unsatisfied with current technology as their predecessors were on the eve of World War II.

Although at any given moment there has always been a computer that was the fastest in the world, the term "supercomputer" began to be commonly used for the fastest computers only with the introduction in 1976 of the Cray-1 supercomputer, manufactured by Cray Research. Fifteen years later, the Cray Y-MP supercomputer from the same company was 16 times faster than the Cray-1. The demand for supercomputers has now brought forth a new generation of companies, each with a "better idea" of how to make machines that are faster yet. By 1996 we can expect supercomputers to reach speeds some 500 times faster than those of only five years earlier. The pace of change is not only unrelenting, it is also drastically accelerating. As the power of supercomputers increases, so will the power of scientists to create and manipulate digital worlds at will, propelling humanity toward new levels of insight and comprehension.

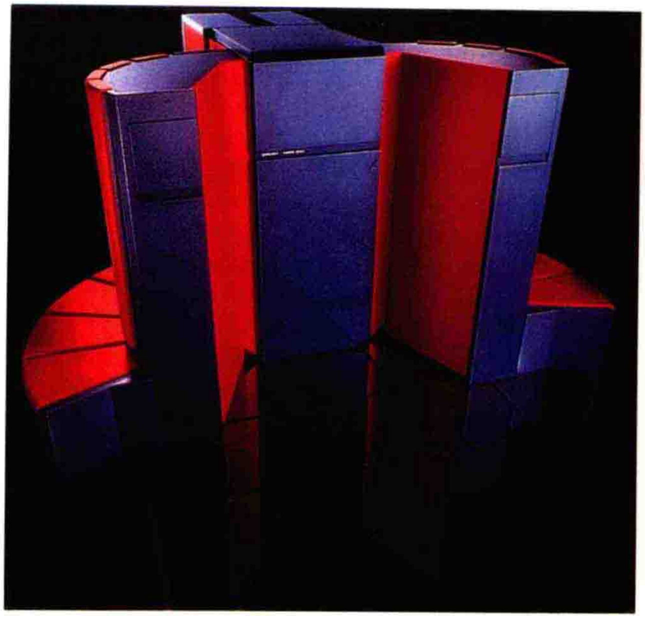

*At peak speed, an eight-processor Cray Y-MP can perform two trillion arithmetic operations per second. The computer's circuit modules are located behind the vertical panels. The cushioned, benchlike arrangement around the base of the computer contains the computer's power supply, as well as some of the plumbing which circulates refrigerating coolant.*

## The Three Modes of Science

For nearly four centuries, science has been progressing primarily through the application of two distinct methodologies: experiment and theory. The experimental/observational mode, first exploited by Galileo in the early 1600s, uses instruments like telescopes, microscopes, and particle accelerators to search for regularities and patterns in the enormous complexity exhibited by natural phenomena. The goal of the experimental branch is to discover facts from which physical models of reality emerge. We refer to such models when we speak today of molecules, viruses, galaxies, and the age of the universe. From William Harvey's discovery of the circulation of blood to Ernest Rutherford's proof that atoms have nuclei, the experimental/observational mode has given us fundamental insights into the world around us.

The theoretical mode, epitomized by the work of Isaac Newton in the mid-1600s, strives to encode the discovered regularities and patterns of the physical world into a set of relationships between mathematical variables. These relationships are expressed by the equations that form the laws of nature. Spectacular successes of the theoretical mode include the Euler and Navier-Stokes equations governing gas and fluid dynamics and Maxwell's equations, which completely describe the behavior of electricity, magnetism, and electromagnetic fields. Another example of a natural law is the Schrödinger equation, which embodies the tenets of quantum mechanics that describe the submicroscopic world of atoms and electrons. Finally, perhaps the grandest example is the Einstein field equations of general relativity, which relate gravity to the curvature of spacetime.

The two traditional modes of science have distinct limitations, however. For the experimenter, nature is sometimes difficult to investigate. Many of the phenomena that scientists would like to observe are too small, or too far away, or too fleeting to yield readily to scientific scrutiny. Theoreticians seeking a solution for a specific instance of a phenomenon traditionally are able to evaluate mathematically only the simplest scenarios. For instance, purely theoretical methods cannot solve exactly the equations that describe the dynamics of a thunderstorm.

The development of digital computers has transformed the pursuit of science because it has given rise to a third methodology: the computational mode. The intent of this mode is to solve numerically the theorist's mathematical models in their full complexity. A simulation that accurately mimics a complex phenomenon contains a wealth of information about that phenomenon. Variables such as temperature, pressure, humidity, and wind velocity are evaluated at thousands of points by the supercomputer as it simulates the development of a storm, for example. Such data, which far exceed anything that could be gained from launching a fleet of weather balloons, reveals intimate details of what is going on in the storm cloud. Furthermore, the computational scientist can compute the collisions of a few atoms just as easily as the collisions of enormous galaxies; the scales of space and time are simply input parameters of the computer program.

## Exploring Solution Space

The mathematical laws of nature that describe the fundamental workings of natural phenomena express universal relationships between such quantities as energy, mass, momentum, temperature, pressure, and density. These equations are powerful because they govern all possible situations, but they do not tell us anything about a particular situation until we "solve" them. For example, Maxwell's equations are universal, *general* statements about electricity and magnetism, but they do not immediately tell us anything about a specific arrangement of a certain set of electric charges that we might want to investigate. If we want to know the electric field in the vicinity of these charges, we must incorporate information about the locations of the charges into Maxwell's equations and then solve them for that particular case. The data appended to the laws of physics to characterize a specific scenario include the boundary conditions that define the spatial extent of the phenomenon to be modeled and the initial conditions that give starting values for the fundamental variables.

When solving the laws of nature, a theoretical scientist traditionally attempts to find an exact solution in symbolic form; such a solution can be written down in terms of known mathematical functions, like sine, cosine, or $x^2$. The relative simplicity of such functions typically requires the scientist to limit the problem under investigation by making some simplifying restrictions such as requiring the geometry of the problem to have special symmetries or the system to be at rest (in equilibrium) and thus unchanging over time.

The complete collection of all solutions covering every conceivable circumstance expressed by a particular law of nature is called a "solution space." Because of the simplifying restrictions required by exact symbolic solutions, these solutions probe only a small region of solution space, and thus tell us little about how nature can behave if, for instance, the geometry of a system is highly asymmetrical or if the system is vigorously dynamic.

A technique called perturbation theory can be used to explore regions of solution space in close proximity to an exact solution. If an actual system differs only slightly (it is "perturbed") from a simpler system for which an exact solution exists, that solution may be expanded mathematically to cover a slightly wider range of circumstances. Thus, for instance, perturbation methods make it possible to investigate a system that has slightly less symmetry or is very near equilibrium. In spite of their great usefulness, these traditional techniques still leave vast areas of solution space unexplored.

By way of an example, consider water waves produced by winds blowing across the ocean. When in equilibrium, the water and air would be separated by a perfectly flat boundary. Obviously, the solution is trivial to write down: the density on each side of the interface is constant (since the two fluids do not mix or propagate waves), and the air and water are moving at a uniform constant velocity with respect to each other, or not moving at all. The familiar small-amplitude waves induced by a gentle breeze are a perturbed solution of this equilibrium state. The solution can be symbolically represented as a nearly flat interface rippled by sine waves propagating with constant velocity.

These waves represent what scientists term a Kelvin-Helmholtz instability, after two nineteenth-century physicists who studied such phenomena. By examining solutions to the equations of fluid dynamics, they discovered that small imperfections in the surface at the interface will grow to form moving waves when two gases or fluids move slowly relative to each other.

But when the gentle breeze increases its velocity to become a raging gale, the simple rolling wave develops much more complex shapes, as the ocean's surface becomes a seething tempest of breaking waves and whitecaps. Here, in a general region of solution space, there exists no exact formula to capture the complexity of the solution representing these phenomena.

## From the Continuous to the Discrete

During the 1940s and 1950s, the Hungarian-American mathematician John von Neumann described a general procedure to explore those regions of solution space beyond the reach of exact solutions. To study complex phenomena, we must leave the continuous world behind and substitute in its place a discrete representation of the phenomena. Then, using a supercomputer, we can indeed solve the relevant equations of fluid flow even if the wind is blowing faster than the speed of sound!

The laws of nature govern what happens at every point in space and at every moment of time. Fortunately, to carry out our science, we do not need to compute solutions to these laws that finely. We can apply our equations at a set of points in space that are separated by distances that are small compared to the size of the objects we care to study. We can study the evolution of the system at discrete intervals of time that are short compared to the duration of the process we are exploring. This replacement of the continuous mathematics by discrete points in space and discrete intervals of time is the key that permits us to use supercomputers to explore the perennially forbidden regions of solution space.

This approach actually predates von Neumann. Consider the archetypical formula for the sine wave: sin ($x$), which can represent the solution to a wave on a plucked string. This formula is very convenient for symbolic computing, as when one needs to manipulate trigonometric identities in order to prove a theorem. However, when we need to use the sine function to solve a practical problem, we evaluate the function at a discrete set of points and exhibit the solution in the familiar table of numbers found in any book on trigonometry, thereby achieving a numerical solution. Indeed one of the first uses of the early mechanical calculators hundreds of years ago was to compute trigonometry and logarithm tables from formulas.

Neither the formula "sin ($x$)" nor the tables of numbers by themselves immediately call to mind the beautiful regularity of this function. To "see" this regularity we graph the numbers, converting the solution from a digital to a visual form that the human brain (some half of whose nerve cells are devoted to visual processing) can easily grasp and associate with physical phenomena, such as rolling waves on the ocean's surface.

This simple methodology of first solving the mathematical laws of nature in a discrete fashion using computers, then converting the numbers to visual images so that the human mind can extract understanding, is the funda-

**SYMBOLIC**

$\sin^2(x) + \cos^2(x) = 1$ $\quad\quad$ $\sin(x + 2\pi) = \sin(x)$

**NUMERIC**

| $x$ | sin ($x$) |
|---|---|
| 0.00 | 0.00 |
| 0.20 | 0.20 |
| 0.40 | 0.39 |
| 0.60 | 0.56 |
| 0.80 | 0.72 |
| 1.00 | 0.84 |
| 1.20 | 0.93 |
| 1.40 | 0.99 |
| 1.60 | 1.00 |
| 1.80 | 0.97 |
| 2.00 | 0.91 |

| $x$ | sin ($x$) |
|---|---|
| 2.20 | 0.81 |
| 2.40 | 0.68 |
| 2.60 | 0.52 |
| 2.80 | 0.33 |
| 3.00 | 0.14 |
| 3.20 | −0.06 |
| 3.40 | −0.26 |
| 3.60 | −0.44 |
| 3.80 | −0.61 |
| 4.00 | −0.76 |
| 4.20 | −0.87 |

| $x$ | sin ($x$) |
|---|---|
| 4.40 | −0.95 |
| 4.60 | −0.99 |
| 4.80 | −1.00 |
| 5.00 | −0.96 |
| 5.20 | −0.88 |
| 5.40 | −0.77 |
| 5.60 | −0.63 |
| 5.80 | −0.46 |
| 6.00 | −0.28 |
| 6.20 | −0.08 |
| 6.40 | 0.12 |

**VISUAL**

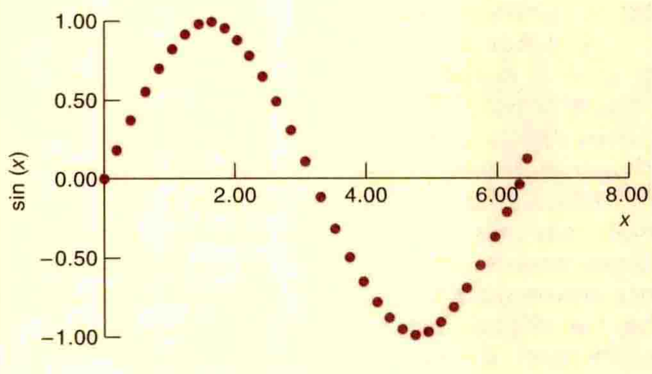


*The simple function, sin (x), can appear as a term in a symbolic statement (top), or it can be evaluated at discrete points to produce a numerical solution (left), which can be converted to graphical form (right).*