Breaking Intractability

Problems that would otherwise be impossible to solve can now be computed, as long as one settles for what happens on the average

by Joseph F. Traub and Henryk Woźniakowski

Although mathematicians and scientists must rank among the most rational people in the world, they will often admit to falling prey to a curse. Called the curse of dimension, it is one many people experience in some form. For example, a family's decision about whether to refinance their mortgage with a 15- or 30-year loan can be extremely difficult to make, because the choice depends on an interplay of monthly expenses, income, future tax and interest rates and other uncertainties. In science, the problems are more esoteric and arguably much harder to cope with. In the computer-aided design of pharmaceuticals, for instance, one might need to know how tightly a drug candidate will bind to a biological receptor. Assuming a typical number of 8,000 atoms in the drug, the biological receptor and the solvent, then because of the three spatial variables needed to describe the position of each atom, the calculation involves 24,000 variables. Simply put, the more variables, or dimensions, there are to consider, the harder it is to accomplish a task. For many problems, the difficulty grows exponentially with the number of variables.

A potentially intractable problem

The curse of dimension can elevate tasks to a level of difficulty at which they become intractable. Even though scientists have computers at their disposal, problems can have so many variables that no future increase in computer speed will make it possible to solve them in a reasonable amount of time.

Can intractable problems be made tractable—that is, solvable in a relatively modest amount of computer time? Sometimes the answer is, happily, yes. But we must be willing to do without a guarantee of achieving a small error in our calculations. By settling for a small error most of the time (rather than always), some kinds of multivariate problems become tractable. One of us (Woźniakowski) formally proved that such an approach works for at least two classes of mathematical problems that arise quite frequently in scientific and engineering tasks. The first is integration, a fundamental component of the calculus. The second is surface reconstruction, in which pieces of information are used to reconstruct an object, a technique that is the basis for medical imaging.

Fields other than science can benefit from ways of breaking intractability. For example, financial institutions often have to assign a value to a pool of mortgages, which is affected by mortgagees who refinance their loans. If we assume a pool of 30-year mortgages and permit refinancing monthly, then this task contains 30 years times 12 months, or 360 variables. Adding to the difficulty is that the value of the pool depends on interest rates over the next 30 years, which are of course unknown.

We shall describe the causes of intractability and discuss the techniques that sometimes allow us to break it. This issue belongs to the new field of information-based complexity, which examines the computational complexity of problems that cannot be solved exactly. We shall also speculate briefly on how information-based complexity might enable us to prove that certain scientific questions can never be answered because the necessary computing resources do not exist in the universe. If so, this condition would set limits on what is scientifically knowable.

Information-based complexity focuses on the computational difficulty of so-called continuous problems. Calculating the movement of the planets is an example. The motion is governed by a system of ordinary differential equations—that is, equations that describe the positions of the planets as a function of time. Because time can take any real value, the mathemati-

cal model is said to be continuous. Continuous problems are distinct from discrete problems, such as difference equations in which time takes only integer values. Difference equations appear in such analyses as the predicted number of predators in a study of predator-prey populations or the anticipated pollution levels in a lake.

In the everyday process of doing science and engineering, however, continuous mathematical formulations predominate. They include a host of problems, such as ordinary and partial differential equations, integral equations, linear and nonlinear optimization, integration and surface reconstruction. These formulations often involve a large number of variables. For example, computations in chemistry, pharmaceutical design and metallurgy often entail calculations of the spatial positions and momenta of thousands of particles.

Often the intrinsic difficulty of guaranteeing an accurate numerical solution grows exponentially with the number of variables, eventually making the problem computationally intractable. The growth is so explosive that in many cases an adequate numerical solution cannot be guaranteed for situations comprising even a modest number of variables.

To state the issue of intractability more precisely and to discuss possible cures, we will consider the example of computing the area under a curve. The process resembles the task of computing the vertical area occupied by a collection of books on a shelf. More explicitly, we will calculate the area between two bookends. Without loss of generality, we can assume the bookends rest at 0 and 1. Mathematically, this summing process is called the computation of the definite integral. (More accurately, the area is occupied by an infinite number of books, each infinitesimally thin.) The mathematical input to this problem is called the integrand, a function that describes the profile of the books on the shelf.

Calculus students learn to compute the definite integral by following a set of prescribed rules. As a result, the students arrive at the exact answer. But most integration problems that arise in practice are far more complicated, and the symbolic process learned in school cannot be carried out. Instead the integral must be approximated numerically—that is, by a computer. More exactly, one computes the integrand values at finitely many points. These integrand values result from so-called information operations. Then one combines these values to produce the answer.

Knowing only these values does not completely identify the true integrand. Because one can evaluate the integrand only at a finite number of points, the information about the integrand is partial. Therefore, the integral can, at best, only be approximated. One typically specifies the accuracy of the approximation by stating that the error of the answer falls within some error threshold. Mathematicians represent this error with the Greek letter epsilon, ε.

Even this goal cannot be achieved without further restriction. Knowing the integrand at, say, 0.2 and 0.5 indicates nothing about the curve between those two points. The curve can assume any shape between them and therefore enclose any area. In our bookshelf analogy, it is as if an art book has been shoved between a run of paperbacks. To guarantee an error of at most ε, some global knowledge of the integrand is needed. One may need to assume, for example, that the slope of the function is always less than 45 degrees—or that only paperbacks are allowed on that shelf.

In summary, an investigator trying to solve an integral must usually do it numerically on a computer. The input to the computer is the integrand values at some points. The computer produces an output that is a number approximating the integral.

The basic concept of computational complexity can now be introduced. We want to find the intrinsic difficulty of solving the integration problem. Assume that determining integrand values and using combinatory operations, such as addition, multiplication and comparison, each have a given cost. The cost could simply be the amount of time a computer needs to perform the operation. Then the computational complexity of this integra-

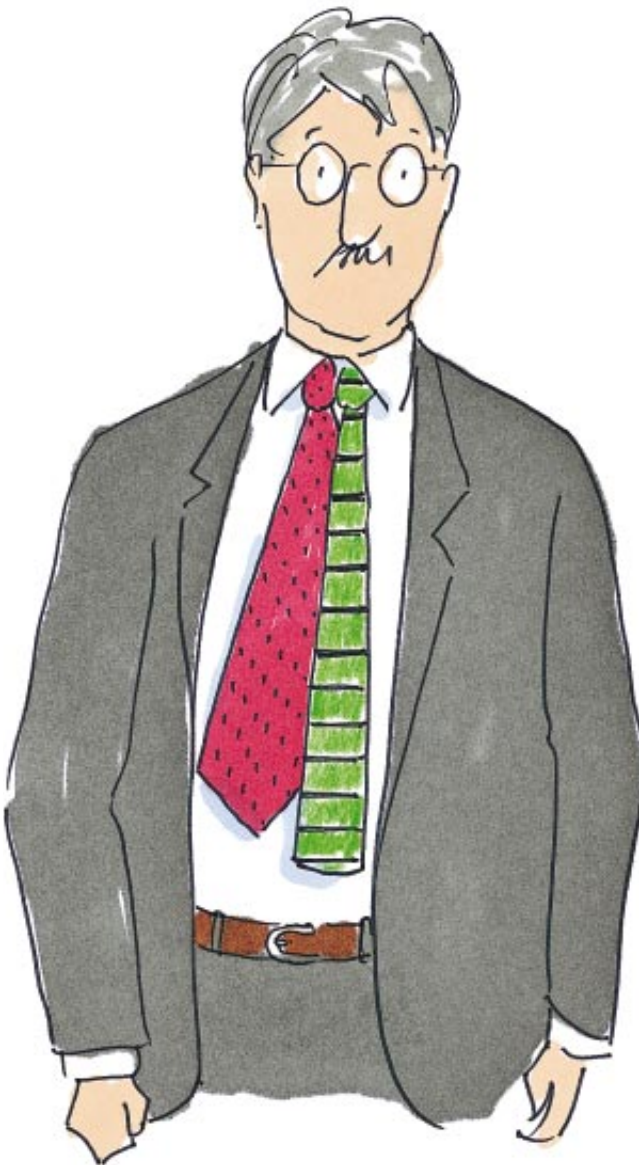

One solution to an intractable problem

JOSEPH F. TRAUB and HENRYK WOŹNIAKOWSKI have been collaborating since 1973. Currently the Edwin Howard Armstrong Professor of Computer Science at Columbia University, Traub headed the computer science department at Carnegie Mellon University and was founding chair of the Computer Science and Telecommunications Board of the National Academy of Sciences. In 1959 he began his pioneering research in what is today called information-based complexity and has received many honors, including election to the National Academy of Engineering. He is grateful to researchers at the Santa Fe Institute for numerous stimulating conversations concerning the limits of scientific knowledge. Woźniakowski holds two tenured appointments, one at the University of Warsaw and the other at Columbia University. He directed the department of mathematics, computer science and mechanics at the University of Warsaw and was the chairman of Solidarity there. In 1988 he received the Mazur Prize from the Polish Mathematical Society. The authors thank the National Science Foundation and the Air Force Office of Scientific Research for their support.

SAMPLING POINTS indicate where to evaluate functions in the randomized and average-case settings. The points are plotted in two dimensions for visual clarity. The points chosen can be spaced over regular intervals such as grid points (*a*), or in random positions (*b*). Two other types, so-called Hammersley points (*c*) and hyperbolic-cross points (*d*), represent optimal places in the average-case setting.

a

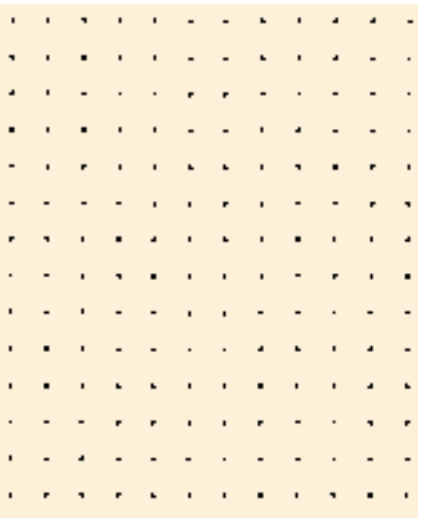

tion problem can be defined as the minimal cost of guaranteeing that the computed answer is within an error threshold, ε, of the true value. The optimal information operations and the optimal combinatory algorithm are those that minimize the cost.

Theorems have shown that the computational complexity of this integration problem is on the order of the reciprocal of the error threshold ($1/\varepsilon$). In other words, it is possible to choose a set of information operations and a combinatory algorithm such that the solution can be approximated at a cost of about $1/\varepsilon$. It is impossible to do better. With one variable, or dimension, the problem is rather easy. The computational complexity is inversely proportional to the desired accuracy.

But if there are more dimensions to this integration problem, then the computational complexity scales exponentially with the number of variables. If d represents the number of variables, then the complexity is on the order of $(1/\varepsilon)^d$—that is, the reciprocal of the error threshold raised to a power equal to the number of variables. If one wants eight-place accuracy (down to 0.00000001) in computing an integral that has three variables, then the complexity is roughly 10^{24}. In other words, it would take a trillion trillion integrand values to achieve that level of accuracy. Even if one generously assumes the existence of a sequential computer that performs 10 billion function evaluations per second, the job would take 100 trillion seconds, or more than three million years. A computer with a million processors would still take 100 million seconds, or about three years.

To discuss multivariate problems more generally, we must introduce one additional parameter, called r. This parameter represents the smoothness of the mathematical inputs. By smoothness, we mean that the inputs consist of functions that do not have any sudden or dramatic changes. (Mathematicians say that all partial derivatives of the function up to order r are bounded.) The parameter takes on nonnegative integer values; increasing values indicate more smoothness. Hence, $r = 0$ represents the least amount of smoothness (technically, the integrands are only continuous—they are rather jagged but still connected as a single curve).

Numerous problems have a computational complexity that is on the order of $(1/\varepsilon)^{d/r}$. For those of a more technical persuasion, multivariate integration, surface reconstruction, partial differential equations, integral equations and nonlinear optimization all have this computational complexity.

If the error threshold and the smoothness parameter are fixed, then the computational complexity depends exponentially on the number of dimensions. Hence, the problems become intractable for high dimensions. An impediment even more serious than intractability may occur: a problem may be unsolvable. A problem is unsolvable if one cannot compute even an approximation at finite cost. This is the case when the mathematical inputs are continuous but jagged. The smoothness pa-

Developing a Random Approach

In the 1940s physicists working on the Manhattan Project at Los Alamos National Laboratory realized that some of the problems they were trying to solve, such as the movement of neutrons through materials, lay beyond the reach of deterministic calculations. They turned to the Monte Carlo method of Nicholas C. Metropolis and Stanislaw M. Ulam. The strength of the method is that its error does not depend on the number of variables in the problem. Hence, if applicable, it breaks the curse of dimension. The classical Monte Carlo method for multivariate integration requires at most of order $1/\varepsilon^2$ evaluations at random points, where ε is the error bound. An alternative statement is that if the integrand is evaluated at n random points, then the expected error of randomization is at most of order $1/\sqrt{n}$. Since its formulation, the Monte Carlo method and its variations have proved to be useful to calculate a variety of phenomena, from the size of cosmic showers to the percolation of a liquid through a solid.

Stanislaw M. Ulam, 1909–84

For multivariate integration, the classical Monte Carlo method is optimal only if the smoothness parameter, r, of integrands is zero. In 1959 the Russian mathematician N. S. Bakhvalov began pioneering research on the computational complexity of multivariate integration in the randomized setting and devised an alternative to the Monte Carlo method. Later, in 1988, Erich Novak of the University of Erlangen-Nürnberg extended the work of Bakhvalov to establish that the computational complexity in the randomized setting is of order $(1/\varepsilon)^s$, with $s = 2/(1 + 2r/d)$. Note that $0 < s \leq 2$. If the smoothness parameter equals zero, then $s = 2$, and the classical Monte Carlo method is optimal. On the other hand, if r is positive, then the classical Monte Carlo method is no longer optimal, and Bakhvalov's method can be used instead.

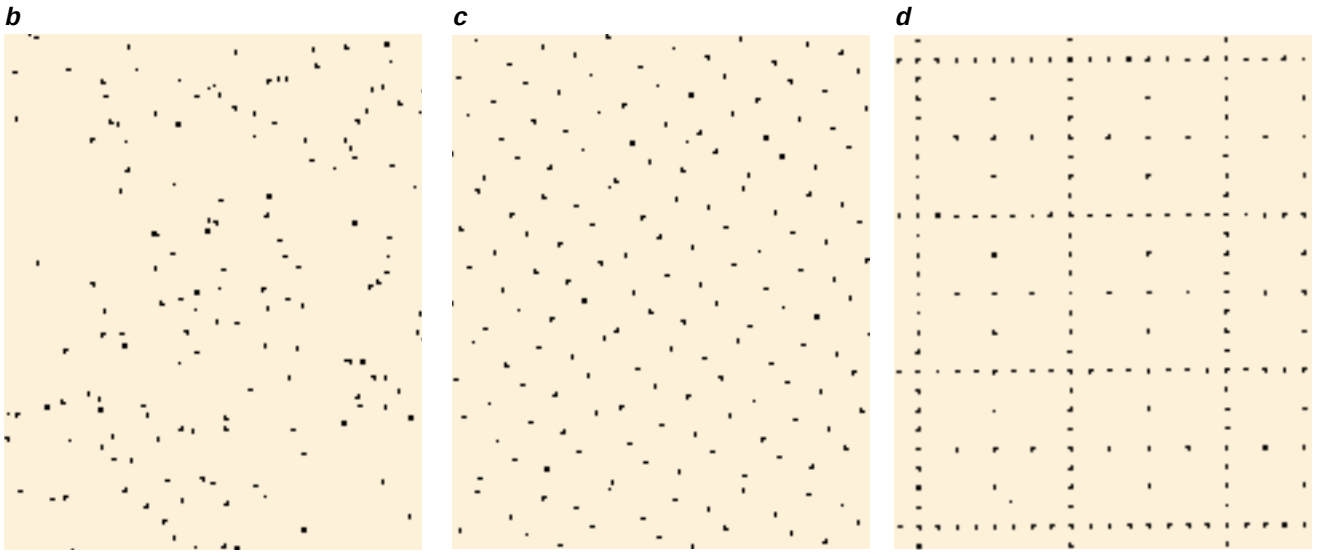

rameter is zero, and the computational complexity becomes infinite. Hence, for many problems with a large number of variables, guaranteeing that an approximation has a desired error becomes an unsolvable or intractable task.

Mathematically, the computational complexity results we have described apply to the so-called worst-case deterministic setting. The "worst case" phrasing comes from the fact that the approximation provides a guarantee that the error always falls within ε. In other words, for multivariate integration, an approximation within the error threshold is guaranteed for every integrand that has a given smoothness. The word "deterministic" arises from the fact that the integrand is evaluated at deterministic (in contrast to random) points.

In this worst-case deterministic setting, many multivariate problems are unsolvable or intractable. Because these results are intrinsic to the problem, one cannot get around them by inventing other methods.

One possible way to break unsolvability and intractability is through randomization. To illustrate how randomization works, we will again use multivariate integration. Instead of picking points deterministically or even optimally, we allow (in an informal sense) a coin toss to make the decisions for us. A loose analogy might be sampling polls. Rather than ask every registered voter, a pollster conducts a small, random sampling to determine the likely winner.

Theorems indicate that with a random selection of points, the computational complexity is at most on the order of the reciprocal of the square of the error threshold ($1/\varepsilon^2$). Thus, the problem is always tractable, even if the smoothness parameter is equal to zero.

The workhorse of the randomized approach has been the Monte Carlo method. Nicholas C. Metropolis and Stanislaw M. Ulam suggested the idea in the 1940s. In the classical Monte Carlo method the integrand is evaluated at uniformly distributed random points. The arithmetic mean of these function values then serves as the approximation of the integral.

Amazingly enough, for multivariate integration problems, randomization of this kind makes the computational complexity independent of dimension. Problems that are unsolvable or intractable if computed from the best possible deterministic points become tractable if approached randomly. (If r is positive, however, then the classical Monte Carlo method is not the optimal one; see box on the opposite page.)

One does not get so much for nothing. The price that must be paid for breaking the unsolvability or intractability is that the ironclad guarantee that the error is at most ε is lost. Instead one is left only with a weaker guarantee that the error is probably no more than ε—much as a preelection poll is usually correct but might, on occasion, predict a wrong winner. In other words, a worst-case guarantee is impossible; one must be content with a weaker assurance.

Randomization makes multivariate integration and many other important problems computationally feasible. It is not, however, a cure-all. Randomization fails completely for some kinds of problems. For instance, in 1987 Greg W. Wasilkowski of the University of Kentucky showed that randomization does not break intractability for surface re-

Average-Case Complexity

In the text, we mention that the average-case complexity of multivariate integration is on the order of the reciprocal of the error threshold ($1/\varepsilon$) and that for surface reconstruction, it is the square of that reciprocal ($1/\varepsilon^2$). For simplicity, we ignored some multiplicative factors that depend on d and ε. Here we provide more rigorous statements.

The average computational complexity, $\mathrm{comp}^{\mathrm{avg}}(\varepsilon, d; \mathrm{INT})$, of multivariate integration is bounded by

$$\frac{g_1(d)}{\varepsilon}\left(\log\frac{1}{\varepsilon}\right)^{(d-1)/2} \le \mathrm{comp}^{\mathrm{avg}}(\varepsilon, d; \mathrm{INT}) \le \frac{g_2(d)}{\varepsilon}\left(\log\frac{1}{\varepsilon}\right)^{(d-1)/2}$$

The average computational complexity, $\mathrm{comp}^{\mathrm{avg}}(\varepsilon, d; \mathrm{SUR})$, of surface reconstruction is bounded by

$$\frac{g_3(d)}{\varepsilon^2}\left(\log\frac{1}{\varepsilon}\right)^{2(d-1)} \le \mathrm{comp}^{\mathrm{avg}}(\varepsilon, d; \mathrm{SUR}) \le \frac{g_4(d)}{\varepsilon^2}\left(\log\frac{1}{\varepsilon}\right)^{2(d-1)}$$

Good estimates of $g_1(d)$, $g_2(d)$, $g_3(d)$ and $g_4(d)$ are currently not known.

construction. Is there an approach that does and that works over a broad class of mathematics problems?

There is indeed. It is the average-case setting, in which we seek to break unsolvability and intractability by replacing a worst-case guarantee with a weaker one: that the expected error is at most ε. The average-case setting imposes restrictions on the kind of mathematical inputs. These restrictions are chosen to represent what would happen most of the time. Technically, the constraints are described by probability distributions; the distributions describe the likelihood that certain inputs occur. The most commonly used distributions are Gaussian measures and, in particular, Wiener measures.

Although it was known since the 1960s that multivariate integration is tractable on the average, the proof was nonconstructive. That is, it did not specify the optimal points to evaluate the integrand, the optimal combinatory algorithm and the average computational complexity. Attempts to apply ideas from other areas of computation to determine these unknowns did not work. For example, evaluating the integrand at regularly spaced points, such as those on a grid, are often used in computation. But theorems have shown them to be poor choices for the average-case setting. A proof was given in 1975 by N. Donald Ylvisaker of the University of California at Los Angeles. It was later generalized in 1990 by Wasilkowski and Anargyros Papageorgiou, then studying for his Ph.D. at Columbia University.

The solution came in 1991, when Woźniakowski found the construction. As sometimes happens in science, a result from number theory, a branch of mathematics far removed from average-case complexity theory, was crucial. Part of the key came from work on number theory by Klaus F. Roth of Imperial College, London, a 1958 Fields Medalist. Another part was provided by recent work by Wasilkowski.

Let us describe the result more precisely. First, put the smoothness parameter at zero—that is, tackle a problem that is unsolvable in the worst-case deterministic setting. Next, assume that integrands are distributed according to a Wiener measure. If we ignore certain multiplicative factors for simplicity's sake, the average computational complexity has been proved to be inversely proportional to the error threshold (on the order of $1/\varepsilon$) [*see box on page 105*]. For small errors, the result is a major improvement over the classical Monte Carlo method, in which the cost is inversely proportional to the square of the error threshold ($1/\varepsilon^2$).

The average case offers a different kind of assurance from that provided by the randomized (Monte Carlo) setting. The error in the average-case setting depends on the distribution of the integrands, whereas the error in the randomized setting depends on a distribution of the sample points. In our books-on-a-shelf analogy, the distribution in the average-case setting might rule out the inclusion of many oversize books, whereas the distribution in the randomized setting determines which books are to be sampled.

In the average-case setting the optimal evaluation points must be deterministically chosen. The best points are Hammersley points or hyperbolic-cross points [*see illustration on pages 104 and 105*]. These deterministic points offer a better sampling than randomly selected or regularly spaced (or grid) points. They make what would be impossible to solve tractable on average.

Is surface reconstruction also tractable on the average? This query is particularly important because, as already mentioned, randomization does not help. Under the same assumptions we used for integration, we find that the average computational complexity is on the order of $1/\varepsilon^2$. Hence, surface reconstruction becomes tractable on average. As was the case for integration, hyperbolic-cross points are optimal.

We are now testing whether the average case is a practical alternative. A Ph.D. student at Columbia, Spassimir H. Paskov, is developing software to compare the deterministic techniques with Monte Carlo methods for integration. Preliminary results obtained by testing certain finance problems suggest the superiority of the deterministic methods in practice.

In our simplified description, we ignored a multiplicative factor that affects the computational complexity. This factor depends on the number of variables in the problem. When the number of variables is large, that factor can become huge. Good theoretical estimates of the factor are not known, and obtaining them is believed to be very hard.

Woźniakowski uncovered a solution: get rid of that factor. Specifically, we say a problem is strongly tractable if the number of function evaluations needed

Discrete Computational Complexity

This article discusses intractability and breaking of intractability for multivariate integration and surface reconstruction. These are two examples of continuous problems. But what is known about the computational complexity of discrete, rather than continuous, problems? The famous traveling salesman problem is an example of a discrete problem, in which the goal is to visit various cities in the shortest distance possible.

A discrete problem is intractable if its computational complexity increases exponentially with the number of its inputs. The intractability of many discrete problems in the worst-case deterministic setting has been conjectured but not yet proved. What is known is that hundreds of discrete problems all have essentially the same computational complexity. That means they are all tractable or all intractable, and the common belief among experts is that they are all intractable. For technical reasons, these problems are said to be NP-complete. One of the great open questions in discrete computational complexity theory is whether the NP-complete problems are indeed intractable [see "Turing Machines," by John E. Hopcroft; SCIENTIFIC AMERICAN, May 1984].

for the solution is completely independent of the number of variables. Instead it would depend only on a power of $1/\varepsilon$. The possibility seems too much to hope for, but it was proved last year that multivariate integration and surface reconstruction are both strongly tractable on the average.

A final obstacle must be overcome before these new results can be used. We know there must exist evaluation points and a combinatory algorithm that make integration and surface reconstruction strongly tractable on the average. Unfortunately, the proof of this result does not tell us what the points and algorithms are, thus leaving a beautiful challenge for the future.

REENTRY OF SPACE SHUTTLE provides an example of a computationally complex task: modeling of the airflow around the craft. This job is difficult even though only seven variables govern the dynamics. Added dimensions may yield problems that can never be solved and thus limit what is scientifically knowable.

Work on information-based complexity has led one of us (Traub) to speculate that it might be possible to prove formally that certain scientific questions are unanswerable. The proposed attack is to prove that the computing resources (time, memory, energy) do not exist in the universe to answer such questions.

One important achievement of mathematics over the past 60 years is the idea that mathematical problems may be undecidable, noncomputable or intractable. Kurt Gödel proved the first of these results. He established that in a sufficiently rich mathematical system, such as arithmetic, there are theorems that can never be proved.

We believe it is time to up the ante and try to prove there are unanswerable scientific questions. In other words, we would like to establish a physical Gödel's theorem. The process offers a markedly different challenge from proving results about mathematical problems, because a scientific question does not come equipped with a mathematical formulation. Such questions include when the universe will stop expanding and what the average global temperature will be in the year 2001.

Why do intractability results suggest that some scientific questions might be unanswerable? Recall the results. In the worst-case deterministic setting, the computational complexity of many continuous problems grows exponentially with dimension. Also, the computational complexity of many discrete problems is conjectured to grow exponentially with the number of inputs [*see box on opposite page*]. Furthermore, although some problems are tractable in the randomized or average-case settings, it has been proved that others remain intractable. Such problems may lurk in certain supercomputing tasks or questions regarding the foundations of physics. After all, they involve a large number of variables or particles. Even worse, many physics problems require solutions to a kind of integral called a path integral, which has an infinite number of dimensions. Solutions of path integrals invite high-dimensional approximations. Thus, the intractability results and conjectures are certainly daunting because they suggest that many tasks with a large number of variables or objects might be impossible to solve.

We emphasize the possibility of other impediments to answering scientific questions. One is chaos, the extreme sensitivity to initial conditions. Because the precise initial conditions are either not known or cannot be exactly entered into a digital computer, certain questions about the behavior of a chaotic system cannot be answered. To focus on the issue at hand, we limit ourselves to intractability.

As we have already observed, a scientific question does not come equipped with a mathematical formulation. Each of a number of models might capture the essence of a scientific question. Because intractability results refer to a particular mathematical formulation, it might happen that although a particular mathematical formulation is intractable, another formulation may be found that is indeed tractable. This prospect indicates a possible way to prove the existence of unanswerable scientific questions. We can attempt to show that there exist scientific questions such that every mathematical formulation that captures the essence of the question is intractable. We would therefore have science's version of Gödel's theorem.

Humans are intrigued not only by the unknown but also by the unknowable. Here we have suggested one possible attack to establish what may be forever unknowable in science. The curse of dimension, broken now for many kinds of problems, may yet cast its spell.

FURTHER READING

INFORMATION-BASED COMPLEXITY. E. W. Packel and J. F. Traub in *Nature,* Vol. 328, No. 6125, pages 29–33; July 2, 1987.

INFORMATION-BASED COMPLEXITY. J. F. Traub, G. W. Wasilkowski and H. Woźniakowski. Academic Press, 1988.

AVERAGE CASE COMPLEXITY OF MULTIVARIATE INTEGRATION. H. Woźniakowski in *Bulletin of the American Mathematical Society,* Vol. 24, No. 1, pages 185–194; January 1991.

THE COMPUTATIONAL COMPLEXITY OF DIFFERENTIAL AND INTEGRAL EQUATIONS: AN INFORMATION-BASED APPROACH. Arthur G. Werschulz. Oxford University Press, 1991.

THEORY AND APPLICATIONS OF INFORMATION-BASED COMPLEXITY. J. F. Traub and H. Woźniakowski in *1990 Lectures in Complex Systems, Santa Fe Institute.* Edited by Lynn Nadel and Daniel L. Stein. Addison-Wesley, 1991.

WHAT IS SCIENTIFICALLY KNOWABLE? J. F. Traub in *Carnegie Mellon University Computer Science: A 25th Anniversary Commemorative.* Edited by Richard F. Rashid. Addison-Wesley, 1991.