# The Great Mathematical Sputnik of 1979

*Eugene L. Lawler*
University of California, Berkeley

## I

Under the frontpage headline, "A Soviet Discovery Rocks World of Mathematics," the New York *Times* of November 7, 1979 announced an event which its readers could easily believe had the importance of the launching of Sputnik. "A surprise discovery by an obscure Soviet mathematician," said the *Times*, "has rocked the world of mathematics and computer analysis ... Apart from its profound theoretical interest, the new discovery may be applicable in weather prediction, complicated industrial processes, petroleum refining, the scheduling of workers at large factories ..." Furthermore, confided the *Times*, "the theory of [secret] codes could eventually be affected by the Russian discovery, and this fact has obvious importance to intelligence agencies everywhere." One could almost hear alarm bells ringing in the offices of the CIA and NSA.

In England, the *Guardian* broke the story three days earlier, under the headline, "Soviet Answer to 'Traveling Salesmen.'" Choosing to emphasize the human side of the story, the *Guardian* said:

> "A young Soviet mathematician, apparently totally unknown to any of the world's senior practitioners, has found an answer to one of the most baffling problems in computer calculation.
>
> "But his obscurity is such that his discovery went unnoticed for 10 months in the mathematical world, although work on the problem has been going on for years.
>
> "The apparent breakthrough was achieved by L. G. Khachian, and was published in a Soviet journal, *Doklady*, last January. Few people in the West read the journal and it was only after rumours of the discovery circulated at a conference in Germany that anyone in the mathematical world at large had even a hint that someone had come up with an answer to what is known in the trade as the 'Traveling Salesman' problem."

So it was that readers of the *Times* and the *Guardian*, and a multitude of other daily newspapers, were told of the Russian mathematical sputnik. But what was this new sputnik? The 1957 sputnik was simply thrown into space, a fact which anyone could comprehend. This 1979 sputnik was, what? A mathematical formula? If a mathematician were to throw it onto a blackboard, what sense could any ordinary person make of it? The *Times* story was so garbled and bereft of technical information, and the *Guardian* was so blatantly wrong, that even professional mathematicians were perplexed and misled.

In this account I shall try to interpret the Russian discovery and place it in proper perspective. At the same time, as the person credited with being the discoverer of the Russian discovery ("Proper attention was finally focussed on the Russian paper by Dr. Eugene Lawler of the University of California at Berkeley.") I am in a favored position to tell how the results of an obscure Russian mathematician found their way from a little-read technical journal in Moscow to the pages of the world's leading English language newspapers. This tale may remind the reader of the children's game of "telephone" in which a story, whispered from ear to ear, progessively changes form until it bears hilariously small similarity to its original version.

## II

My own role in this story is a bit less than heroic, and easily explained. The conference referred to by the *Guardian* was held at Oberwolfach, a beautiful Black Forest village near the Swiss border. In this idyllic setting, the Germans maintain a conference center which has been a focal point for the revitalization of German mathematics since World War II. In May, 1979, I was one of a handful of specialists in operations research, econometrics, computer science, and applied mathematics gathered together to discuss "mathematical programming."

A large part of the important business of any scientific meeting is conducted outside formal sessions. Meetings at Oberwolfach are designed to promote informal interchanges, with comfortable lounges well appointed with refrigerators of beer and Rhine wine. At one bull session, Professor Rainer Burkhard of Cologne produced a very blurry Xerox copy of a Russian paper that had been mailed to him a few days before by a friend at the Polish Academy of Sciences in Warsaw. Burkard asked if anyone there knew anything about the paper. I confess that I cannot remember his exact words on this momentous occasion, but I think they may have been something like, "Does anybody here know anything about this?"

No one in our group knew anything about the paper. But one of our number could read Russian. He gave us the title, "Polynomial Algorithm for Linear Programming," and bits and pieces of the text. There was no doubt in any of our minds as to what the paper was about. There was considerable doubt as to whether the Russian's results were correct. And the paper gave only theorems, with no proofs.

Why did we doubt the correctness of the Russian's results? For one thing, we knew that in 1973 a Brazilian mathematician announced a solution to the same problem at an international symposium held at Stanford University. His lecture was mobbed, but scarcely a soul in his audience could make heads nor tails of what he was saying. In the aftermath, wagers were made as to whether or not the Brazilian's proofs would stand up under close scrutiny. I still covet a check for one dollar I won from one prominent theoretician for taking the short side of a ten to one bet.

There were similar wagers made at Oberwolfach. But no one in our group appeared ready to volunteer to spend the time and effort required to check the Russian's claims and to supply proofs. All of us had our own research and, quite honestly, few of us felt completely comfortable with the mathematics used in the paper.

After the conference, I returned to Amsterdam where I was in residence at the time. As luck would have it, a Czech visitor, Milan Vlach, could read Russian. Together we knocked out a loose translation of the paper. I sent copies to a couple of dozen people in the U.S., with hopes that at least one of them would be willing to expend the effort required to give an informed opinion. What I was trying to do, of course, was to get someone else to do the dog work, while I got ahead on my own research.

I had no way of knowing whether or not anyone in the U.S. had already seen the Russian paper. The *Doklady* in which it appeared is hardly obscure: "Doklady" simply means "reports" and the reports in question were the Doklady of the Academy of Sciences of the U.S.S.R. Surface mail from Russia being what it is, the journal actually arrived at the Mathematics Library in Berkeley a few days after I saw the paper at Oberwolfach. Thus, even the most dedicated Russian journal watcher in the U.S. could not have discovered Khachian's paper before I saw it.

When I returned to Berkeley in July, I still had no response to my inquiries. However, two Hungarian mathematicians, Peter Gacs of the University of Rochester, and Laslo Lovasz of the University of Szeged, happened to be visiting at neighboring Stanford that month. With a few days work, they managed to check the Russian paper and supply proofs. The following month in Montreal, they announced their validation of the paper at the 1979 edition of the same international symposium at which the false alarm had been sounded six years before.

III

The public unveiling of Khachian's results occurred with the October 4 cover story of *Science News:* "Linear Programming: Solid New Algorithm." This piece began with a flowery prologue describing abstract mathematics as a "sort of dream." But despite some inaccurate and misleading illustrations, the account was generally correct, and did not seriously mistate the significance of the achievement.

Prompted by the *Science News* story, Gina B. Kolata began preparing an article for *Science* magazine. Kolata researched her article well, talking to a number of mathematicians for background information, including Ronald Graham at Bell Laboratories, Lovasz, and myself. The article was generally accurate and well balanced. One can only speculate as to how the article came to be seriously misinterpreted. However, it did contain one somewhat mushy paragraph in which the words "Khachian's result ... is tied to what is said to be the major unsolved problem in computer science ... the traveling salesman problem." (Here my ellipses are purposely chosen to radically change the meaning.) Most significantly, the article appeared under the title, "Mathematicians Amazed by Russian's Discovery."

Newspaper science writers are more assiduous readers of *Science* than mathematicians are of *Doklady*. And when an authoritative publication like *Science* says that mathematicians are "amazed" by a Russian discovery, any newsman can smell a story. What happened next was almost as inevitable as a lead paragraph following a science writer's byline. The *Guardian* story, communicated by its Washington correspondent, appears to have been based on nothing more than a very careless misreading of the *Science* story. The London *Telegraph* rewrite offered the observation that there are few traveling salesmen in Communist countries. (A fact which the Chicago *Tribune* found worthy of editorial comment.)

The *Times* story appears to have been based on certain unshakable preconceptions of its writer, Malcolm W. Browne. Browne called George Dantzig of Stanford University, a great pioneering authority on linear programming, and tried to force him into various admissions. Dantzig's version of the interview bears repeating.

> "What about the traveling salesman problem?" asked Browne. "If there is a connection, I don't know what it is," said Dantzig. ("The Russian discovery proposed an approach for [solving] a class of problems related to the 'Traveling Salesman Problem,'" reported Browne.) "What about cryptography?" asked Browne. "If there is a connection, I don't know what it is," said

Dantzig. ("The theory of codes could eventually be affected," reported Browne.) "Is the Russian method practical?" asked Browne. "No," said Dantzig. ("Mathematicians describe the discovery ... as a method by which computers can find solutions to a class of very hard problems that has hitherto been attacked on a hit-or-miss basis," reported Browne.)

*Science* is mailed to the 130,000 members of the American Association for the Advancement of Science. Immediately upon publication of its story, my phone began ringing, as did those of Dantzig, Gacs, Graham, and almost everyone else mentioned in the story. Lovasz, back in Hungary, was protected from this onslaught, as of course was Khachian, still in relative isolation in Moscow. When the *Times* story appeared, the deluge predictably increased, and predictably changed character. One caller asked me to explain how in the world the Russians allowed such an important secret to escape their borders. A Beverly Hills lawyer asked if this new development might affect the immigration status of a client. Graham was called by a radio talk show host and asked if it wasn't true that the Russian discovery would change the lives of every man, woman, and child in America.

When *Time* magazine called, I was apprehensive. A cover story? Perhaps Khachian was to be Man of the Year? But no. It was not hard to convince the *Time* man that the story was, well, just a bit overblown and perhaps *Time* should pass it up. His parting question was to ask if I could recommend a good book on linear programming.

IV

What had Khachian done? He had proposed a remarkable new algorithm, or computational procedure, for solving linear programming problems. And, as every modern M.B.A. knows, linear programming has innumerable applications in economic modeling and business planning. It is used to allocate resources, plan production, schedule workers, plan investment portfolios, formulate marketing (and military) strategies. The versatility and economic impact of linear programming in today's industrial world is truly awesome.

In order to understand how large corporations use linear programming to minimize costs or maximize profits, let's consider an application for consumers: buying food at the supermarket. We want our food purchases to cost as little as possible, but we insist that they provide us with all the nutrients required to sustain healthful life.

We must first decide what our daily nutritional requirements are: so many calories, so much protein, fat, and carbohydrate, so many units of each of several vitamins, minerals, and so forth. Let us suppose that we decide upon 30 such nutritional components, and our daily requirement of each of them. We then go to the supermarket in which there are (say) 1,000 different foods to choose from. For each of the foods we record its price per unit and the amount per unit of our daily requirement of each of the 30 different nutritional components which the food supplies. When we are done, we have a table of 31,000 numbers.

Our table of numbers provides the input data for a linear programming problem with 1,000 "decision variables" indicating the amounts of each of the foods to be purchased, and 30 "constraints" determined by the nutritional requirements. Each constraint is linear, in that it involves simply a summation of the 1,000 variables, appropriately weighted by coefficients specified by our table of numbers. Our problem also has a linear "objective function," determined by a summation of the 1,000 variables, with coefficients specified by the prices of the foods. An "optimal" solution to the linear programming problem minimizes the objective function, i.e., makes our purchases of food as cheap as possible, subject to the nutritional constraints. The resulting diet may not be tasteful, consisting of (say) peanut butter sandwiches, mustard greens and buttermilk. But there is no doubt that it will be "rational" and "optimal," within the limitations of the model we have formulated.

A linear programming problem like ours, with 1,000 variables and 30 constraints, is by no means large by present day standards. Problems of much larger size and complexity are routinely formulated and solved, quite successfully, by the "simplex" method of computation invented by George Dantzig.

The simplex method can be visualized in geometric terms. Imagine a space whose dimensionality is determined by the number of variables, a dimensionality of 1,000 in the case of our diet problem. (There is really no magic in working in high dimensional spaces. Mathematicians don't try very hard to visualize them, and usually content themselves with drawing suggestive sketches of three, or even two, dimensional examples.) Within this space imagine a polyhedron (a multidimensional polygon) whose flat sides or "faces" are determined by the linear constraints of the problem and whose corner points or "vertices" correspond to possible solutions. The simplex computation proceeds from one vertex of this polyhedron to another, continually improving the value of the solution, until finally it arrives at a vertex corresponding to an optimal solution.

The computational method that Khachian proposed is quite different, but can also be described in terms of the geometry of the same high dimensional space. At a given point in the computation, an optimal solution is known to lie within an ellipsoid (a multidimensional ellipse, sort of egg shaped). The center of this ellipsoid is tested to see if it is an optimal solution. If it is not, it is possible to slice the ellipsoid in two with a plane defined by one of the linear constraints of the problem. An optimal solution is known to lie within one of the "semiellipsoids" which result from this slicing. This semiellipsoid is then artfully surrounded by a new ellipsoid, and the process repeated. The ellipsoids get smaller and smaller, and eventually the center of one of them turns out to yield an optimal solution.

## V

In order to understand why the new algorithm is so remarkable, we must understand how applied mathematicians and computer scientists measure computational efficiency.

One accepted practice is to determine the number of computational steps (and thus the running time) which an algorithm will require, in the worst possible case, for a problem of a given size. If the worst-case running time increases no faster than a polynomial function of problem size, an algorithm is said to be "polynomial bounded." (A polynomial function of problem size is, for example, the square, cube, or any fixed power of problem size.)

It is a mathematical fact that a polynomial algorithm can be guaranteed to outperform (in the worst case) a nonpolynomial algorithm, for sufficiently large problems. It is also an empirical fact that polynomial algorithms tend to be efficient in practice, and usually outperform nonpolynomial algorithms. Beginning about fifteen years ago, both theoretical and practical computer scientists began to accept that polynomial algorithms are, virtually by definition, "good" and "efficient" algorithms.

Suppose we concern ourselves only with linear programming problems in which the coefficients are whole numbers and not fractions, and suppose we measure problem size by the total number of digits in all the coefficients, e.g., by the total number of digits in the table of 31,000 numbers for the diet problem. Subject to these ground rules, the simplex method is not polynomial bounded. In 1967, Victor Klee of the University of Washington and George Minty of Indiana, concocted a "pathological" class of linear programming problems for which the most commonly employed version of the simplex method must carry out a seemingly endless sequence of computational steps before arriving at an optimal solution.

The ellipsoidal method is polynomial bounded. But does this mean that it is better than the simplext method? The *worst case* behavior of the ellipsoidal method is better, but its average or typical running time is certainly a good deal worse. Dantzig has estimated that problems which are routinely solved within half an hour by the simplex method would require fifty million years to solve by the new method. Some preliminary computational experiments have confirmed these pessimistic projections, and have also indicated some technical difficulties, such as "numerical instability." This latter defect can be overcome by performing extremely precise arithmetic on numbers with (say) hundreds of thousands of digits to the right of the decimal point. But this is a cure which is almost worse than the disease itself.

It is possible that the new algorithm can be modified and improved so that it can become a practical computational method. But this is a matter for much further investigation. And by the time that happens, it is also entirely possible that someone will have found a way to modify the simplex method so that it, too, will be polynomial bounded.

## VI

In order to understand what the *Times* called the "profound theoretical interest" of Khachian's result, its alleged relation to the traveling salesman problem, and how "the theory of codes could eventually be affected," a bit more background information is required. It is an irresistible temptation for me to begin by quoting the opening paragraphs of my own textbook, **Combinatorial Optimization:**

> "Combinatorial analysis is the mathematical study of the arrangement, grouping, ordering, or selection of discrete objects, usually finite in number. Traditionally, combinatorialists have been concerned with questions of existence or of enumeration. That is, does a particular type of arrangement exist? Or how many such arrangements are there?
>
> "Quite recently, a new line of combinatorial investigation has gained increasing importance. The question asked is not 'Does the arrangement exist?' or 'How many arrangments are there?', but rather 'What is a *best* arrangement?' The existence of a particular type of arrangement is usually not in question, and the number of such possible arrangements is irrelevant. All that matters is finding an optimal arrangement, whether it be one in a hundred or one in an effectively infinite number of possibilities."

Specialists in combinatorial optimization have a large stable of pet problems whose colorful names tend to belie their considerable technological and economic importance. There are the "Chinese postman's problem," the "knapsack problem," and even the "homosexual marriage problem." An enduring favorite is the traveling salesman problem, possibly because of its deceptively simple statement: A salesman must visit each of several specified cities. How can he find, with reference to a highway map or mileage table, a tour which will enable him to visit each city exactly once and to return to his home city with the smallest total mileage on the odometer of his car?

In recent years, theoretical computer scientists have become fascinated with issues of computational complexity: Which problems are easy to solve? Which problems are inherently dificult? And why? They have pored over the stable of problems in combinatorial optimization, analyzing them and classifying them, in much the same way that physicists try to bring order into the collection of particles in their subatomic zoo.

A very useful approach to the classification of problems according to their inherent computational complexity was provided in the early 1970's by theoretical results of Stephen Cook of Toronto, coupled with insights of Richard Karp of Berkeley. (Essentially the same results were obtained independently by a Russian, L. A. Levin, now at M.I.T.) Under this approach, problems which can be solved in polynomial bounded time are placed in a class called P. Another class of problems called NP (for "nondeterministic polynomial") is defined. Roughly speaking, the problems in NP are those for which it is possible to prove the correctness of a solution in polynomial bounded time, if one is able to guess what the solution is. Every problem in P is also in NP.

It has not been proved that there are problems in NP which are not in P, but it is known that there are very special problems in NP called NP-complete problems. If any NP-complete problem can be solved in polynomial bounded time then *all* problems in NP can be solved in polynomial bounded time. A specially formulated version of the traveling salesman problem is known to be NP-complete. Hence a polynomial bounded solution to the traveling salesman problem would imply that P=NP.

The "P=NP" question can reasonably be said to be the major open problem in theoretical computer science today. It is a truly baffling problem, one which many investigators have tackled and come away from essentially empty handed. It has even attracted the attention of mathematical logicians, some of whom have speculated on the relations between the P=NP question and the axiomatic foundations of mathematics itself.

Although no one has been able to prove that P is not equal to NP, nearly all computer scientists conjecture that this is the case, and a certain amount of circumstantial evidence has accumulated to support this conjecture. Much current research, including some research in cryptography, proceeds on the premise that certain problems in NP are difficult to solve in practice.

The problem of cryptography is to encode messages so that it will be exceedingly difficult and time consuming for an outsider to unscramble the encoded text. A number of novel proposals for "public key" encryption systems have been made by computer scientists. One well publicized proposal is based on the assumption that factorization of large numbers, a problem in NP, is very difficult. Others are based on the assumption that still other problems in NP are hard.

## VII

Khachian emphatically did not discover a polynomial bounded algorithm for the traveling salesman problem. Had he done so, as the *Guardian* claimed and the *Times* strongly implied, the effect would be devastating. With P proved equal to NP, a possibility discounted by nearly all theoreticians, innumerable technical papers instantly would become dead letters. The theory of computational complexity would require top-to-bottom rethinking. One of the lesser issues would be that various proposals for data encryption would appear ill advised. And if the polynomial bounded algorithm proved to be efficient in a practical sense, its economic impact quite possibly would exceed that of the simplex method.

What Khachian did do was to answer a much smaller question in complexity theory. Linear programming had been known to be a problem in NP. It had not been shown to be NP-complete, and there was some evidence that it could not be. Some investigators, including myself, thought that perhaps linear programming was a problem of intermediate difficulty: not in P, but also not NP-complete. Khachian squeezed linear programming into P and thereby resolved the issue.

But there is even some question as to how much recognition should be given to Khachian for this smaller result. It appears that the ellipsoidal algorithm was developed by three other Russians, D. B. Judin, A. Z. Nemirovsky, and N. Z. Shor, in the context of more general problems than linear programming. In a 1976 paper, Judin and Nemirovsky explicitly suggested the application of the ellipsoidal approach to linear programming. In a 1977 paper (not referenced by Khachian), Shor worked out almost the precise formulas employed by Khachian. What Khachian did was to supply the final bit of argumentation necessary to obtain the polynomial boundedness result. Because of this mixture of contributions (and in deference to the fact that Khachian is of Armenian, rather than Russian, extraction), perhaps the ellipsoidal algorithm should be referred to as the "Soviet" method.

## VIII

On November 11, under the head, "Soviet Mathematician is Obscure No More," the *Times* reported that it had located Khachian. The *Times* Moscow correspondent found that he was "a relaxed, friendly young man," a "kandidat" who performs research at the computer center of the Academy of Sciences. Khachian was reported as saying that he was "somewhat surprised" by the enthusiastic response his paper had had in the West. "All this glory has fallen on him quite unexpectedly," his "teacher and mentor," G. S. Pospelov, was quoted as saying.

One might say that Khachian and Pospelov just don't understand Western journalism as well as they understand mathematics.

In fact, for mathematicians who might have been unfamiliar with the inner workings of the press, the behavior of the *Times* provided a kind of higher education. After Malcolm Browne's first article appeared on November 7, mathematicians telephoned the *Times*, wrote letters to the editors and reporters involved, and invited the *Times* to seminars on the ellipsoid method at IBM and in New York City. But the paper printed three follow-up articles, the last of which repeated the mistake of the first, saying, "Mr. Khachian's method is believed to offer an approach for the linear programming of computers to solve so-called 'traveling salesman' problems."

Only much later, on March 21, did the *Times* print a retraction. And the retraction was so artfully constructed that the uninformed might have mistaken it for a simple update. Six paragraphs down, the piece did admit that the *Times* had been mistaken in its linking of Khachian's work to the traveling salesman problem. But the headline read "A Russian's Solution in Math Questioned" and the subhead read "Americans Who Studied Khachian Linear Programming Method Express Doubt on Scope." That made it sound as though poor Khachian had exaggerated the importance of his work, and Western mathematicians had cut him down to size. The lead sentence of the piece reinforced that impression: "American mathematicians who have studied the new Soviet Method for solving a difficult class of computational problems known as linear programming problems say that the feat announced last November, while important, is far from the seminal achievement originally portrayed."

To which one can only nod, and add with a sigh, "Originally portrayed ... by the *newspapers*."