

Linear Programming: Chapter 16

Interior-Point Methods

Robert J. Vanderbei

November 6, 2007

Operations Research and Financial Engineering
Princeton University
Princeton, NJ 08544

<http://www.princeton.edu/~rvdb>

Interior-Point Methods—The Breakthrough

Breakthrough in Problem Solving

By JAMES GLEICK

A 28-year-old mathematician at A.T.&T. Bell Laboratories has made a startling theoretical breakthrough in the solving of systems of equations that often grow too vast and complex for the most powerful computers.

The discovery, which is to be formally published next month, is already circulating rapidly through the mathematical world. It has also set off a deluge of inquiries from brokerage houses, oil companies and airlines, industries with millions of dollars at stake in problems known as linear programming.

Faster Solutions Seen

These problems are fiendishly complicated systems, often with thousands of variables. They arise in a variety of commercial and government applications, ranging from allocating time on a communications satellite to routing millions of telephone calls over long distances, or whenever a limited, expensive resource must be spread most efficiently among competing users. And investment companies use them in creating portfolios with the best mix of stocks and bonds.

The Bell Labs mathematician, Dr. Narendra Karmarkar, has devised a radically new procedure that may speed the routine handling of such problems by businesses and Government agencies and also make it possible to tackle problems that are now far out of reach.

"This is a path-breaking result," said Dr. Ronald L. Graham, director of mathematical sciences for Bell Labs in Murray Hill, N.J.

"Science has its moments of great progress, and this may well be one of them."

Because problems in linear programming can have billions or more possible answers, even high-speed computers cannot check every one. So computers must use a special procedure, an algorithm, to examine as few answers as possible before finding the best one — typically the one that minimizes cost or maximizes efficiency.

A procedure devised in 1947, the simplex method, is now used for such problems,

Continued on Page A19, Column 1



Karmarkar at Bell Labs: an equation to find a new way through the maze

Folding the Perfect Corner

A young Bell scientist makes a major math breakthrough

Every day 1,200 American Airlines jets crisscross the U.S., Mexico, Canada and the Caribbean, stopping in 110 cities and bearing over 80,000 passengers. More than 4,000 pilots, copilots, flight personnel, maintenance workers and baggage carriers are shuffled among the flights; a total of 3.6 million gal. of high-octane fuel is burned. Nuts, bolts, altimeters, landing gears and the like must be checked at each destination. And while performing these scheduling gymnastics, the company must keep a close eye on costs, projected revenue and profits.

Like American Airlines, thousands of companies must routinely untangle the myriad variables that complicate the efficient distribution of their resources. Solving such monstrous problems requires the use of an abstruse branch of mathematics known as linear programming. It is the kind of math that has frustrated theoreticians for years, and even the fastest and most powerful computers have had great difficulty juggling the bits and pieces of data. Now Narendra Karmarkar, a 28-year-old

Indian-born mathematician at Bell Laboratories in Murray Hill, N.J., after only a year's work has cracked the puzzle of linear programming by devising a new algorithm, a step-by-step mathematical formula. He has translated the procedure into a program that should allow computers to track a greater combination of tasks than ever before and in a fraction of the time.

Unlike most advances in theoretical mathematics, Karmarkar's work will have an immediate and major impact on the real world. "Breakthrough is one of the most abused words in science," says Ronald Graham, director of mathematical sciences at Bell Labs. "But this is one situation where it is truly appropriate."

Before the Karmarkar method, linear equations could be solved only in a cumbersome fashion, ironically known as the simplex method, devised by Mathematician George Dantzig in 1947. Problems are conceived of as giant geodesic domes with thousands of sides. Each corner of a facet on the dome

The Wall Street Journal Waits 'Till 1986

THE STARTLING DISCOVERY BELL LABS KEPT IN THE SHADOWS

Now its breakthrough mathematical formula could save business millions

It happens all too often in science. An obscure researcher announces a stunning breakthrough and achieves instant fame. But when other scientists try to repeat his results, they fail. Fame quickly turns to notoriety, and eventually the episode is all but forgotten.

That seemed to be the case with Narendra K. Karmarkar, a young scientist at AT&T Bell Laboratories. In late 1984 the 28-year-old researcher astounded not only the scientific community but also the business world. He claimed he had cracked one of the thorniest aspects of computer-aided problem-solving. If so, his feat would have meant an instant windfall for many big companies. It could also have pointed to better software for small companies that use computers to help manage their business.

Karmarkar said he had discovered a quick way to solve problems so hideously complicated that they often defy even the most powerful supercomputers. Such problems befall a broad range of business activities, from assessing risk factors in stock portfolios to drawing up production schedules in factories. Just about any company that distributes products through more than a handful of warehouses bumps into such problems when calculating the cheapest routes for getting goods to customers. Even when the problems aren't terribly complex, solving them can chew up so much computer time that the answer is useless before it's found.

HEAD START. To most mathematicians, Karmarkar's precocious feat was hard to swallow. Because such questions are so common, a special branch of mathematics called

twist. Other scientists weren't able to duplicate Karmarkar's work, it turns out, because his employer wanted it that way. Vital details about how best to translate the algorithm, whose mathematical notations run on for about 20 printed pages, into digital computer code were withheld to give Bell Labs a head start at developing commercial products. Following the breakup of American Telephone & Telegraph Co. in January, 1984, Bell Labs was no longer prevented from exploiting its research for profit. While the underlying concept could not be patented or copyrighted because it is pure knowledge, any computer programs that AT&T developed to implement the procedure can be protected.

Now, AT&T may soon be selling the first product based on Karmarkar's work—to the U.S. Air Force. It includes a multiprocessor computer from Alliant Computer Systems Corp. and a software version of Karmarkar's algorithm that has been optimized for high-speed parallel processing. The system would be installed at St. Louis' Scott Air Force Base, headquarters of the Military Airlift Command (MAC). Neither party will comment on the deal's cost or where the negotiations stand, but the Air Force's interest is easy to fathom.

JUGGLING ACT. On a typical day thousands of planes ferry cargo and passengers among air fields scattered around the world. To keep those jets flying, MAC



KARMARKAR: SKEPTICS ATTACKED HIS PRECOCIOUS FEAT

linear programming (LP) has evolved, and most scientists thought that was as far as they could go. Sure enough, when other researchers independently tried to test Karmarkar's process, their results were disappointing. At scientific conferences skeptics attacked the algorithm's validity as well as Karmarkar's veracity.

But this story may end with a different

Karmarkar Algorithm Proves Its Worth

Less than two years after discovery of a mathematical procedure that Bell Labs said could solve a broad range of complex business problems 50 to 100 times faster than current methods, AT&T is filing for patents covering its use. The Karmarkar algorithm, which drew headlines when discovered by researcher Narendra Karmarkar, will be applied first to AT&T's long-distance network.

Thus far, Bell Labs has verified the procedure's capabilities in developing plans for new fiber-optic transmission and satellite capacity linking 20 countries bordering the Pacific Ocean. That jointly owned network will be built during the next 10 years. Planning requires a tremendous number of "what if" scenarios involving 43,000 variables describing transmission capacity, location and construction schedules, all juggled amid political considerations of each connected country.

The Karmarkar algorithm was able to solve the Pacific Basin problem in four minutes, against 80 minutes by the method previously used, says Neil Dinn, head of Bell Labs' international transmission planning department. The speedier solutions will enable international committees to agree on network designs at one meeting instead of many meetings stretched out over months.

AT&T now is using the Karmarkar procedure to plan construction for its domestic network, a problem involving 800,000 variables. In addition, the procedure may be written into software controlling routing of domestic phone calls, boosting the capacity of AT&T's current network.

AT&T Patents the Algorithm, & Announces KORBX

Patents

by Stacy V. Jones

A Method to Improve Resource Allocation

Scientists at Bell Laboratories in Murray Hill, N.J., were granted three patents this week for methods of improving the efficiency of allocation of industrial and commercial resources.

The American Telephone and Telegraph Company, the laboratory's sponsor, is using the methods internally to regulate such operations as long-distance services.

Narendra K. Karmarkar of the laboratory staff was granted patent 4,744,028 for methods of allocating telecommunication and other resources. With David A. Bayer and Jeffrey C. Lagarian as co-inventors, he was granted patent 4,744,027 on improvements of the basic method. Patent 4,744,026 went to Robert J. Vanderbei for enhanced procedures.



Narendra K. Karmarkar of the Bell Laboratories staff.

THE NEW YORK TIMES, May 14, 1988

AT&T Markets Problem Solver, Based On Math Whiz's Find, for \$8.9 Million

By ROGER LOWENSTEIN

Staff Reporter of THE WALL STREET JOURNAL.
NEW YORK—American Telephone & Telegraph Co. has called its math whiz, Narendra Karmarkar, a latter-day Isaac Newton. Now, it will see if he can make the firm some money.

Four years after AT&T announced an "astonishing" discovery by the Indian-born Mr. Karmarkar, it is marketing an \$8.9 million problem solver based on his invention.

Dubbed Korbx, the computer-based system is designed to solve major operational problems of both business and government. AT&T predicts "substantial" sales for the product, but outsiders say the price is high and point out that its commercial viability is unproven.

"At \$9 million a system, you're going to have a small number of users," says Thomas Magnanti, an operations-research specialist at Massachusetts Institute of Technology. "But for very large-scale problems, it might make the difference."

Korbx uses a unique algorithm, or step-by-step procedure, invented by Mr. Karmarkar, a 32-year old, an AT&T Bell Laboratories mathematician.

"It's designed to solve extremely difficult or previously unsolvable resource-allocation problems—which can involve hundreds of thousands of variables—such as personnel planning, vendor selection, and equipment scheduling," says Aristides Fronistas, president of an AT&T division created to market Korbx.

Potential customers might include an airline trying to determine how to route many planes between numerous cities and an oil company figuring how to feed different grades of crude oil into various refineries and have the best blend of refined products emerge.

AT&T says that fewer than 10 companies, which it won't name, are already using Korbx. It adds that, because of the price, it is targeting

only very large companies—mostly in the Fortune 100.

Korbx "won't have a significant bottom-line impact initially" for AT&T, though it might in the long term, says Charles Nichols, an analyst with Bear, Stearns & Co. "They will have to expose it to users and demonstrate" it uses.

AMR Corp.'s American Airlines says it's considering buying AT&T's system. Like other airlines, the Fort Worth, Texas, carrier has the complex task of scheduling pilots, crews and flight attendants on thousands of flights every month.

Thomas M. Cook, head of operations research at American, says, "Every airline has programs that do this. The question is: Can AT&T do it better and faster? The jury is still out."

The U.S. Air Force says it is considering using the system at the Scott Air Force Base in Illinois.

One reason for the uncertainty is that AT&T has, for reasons of commercial secrecy, deliberately kept the specifics of Mr. Karmarkar's algorithm under wraps.

"I don't know the details of their system," says Eugene Bryan, president of Decision Dynamics Inc., a Portland, Ore., consulting firm that specializes in linear programming, a mathematical technique that employs a series of equations using many variables to find the most efficient way of allocating resources.

Mr. Bryan says, though, that if the Karmarkar system works, it would be extremely useful. "For every dollar you spend on optimization," he says, "you usually get them back many-fold."

AT&T has used the system in-house to help design equipment and routes on its Pacific Basin system, which involves 22 countries. It's also being used to plan AT&T's evolving domestic network, a problem involving some 800,000 variables.

THE WALL STREET JOURNAL, August 15, 1988

What Makes LP Hard?

Primal

$$\begin{array}{ll} \text{maximize} & c^T x \\ \text{subject to} & Ax + w = b \\ & x, w \geq 0 \end{array}$$

Dual

$$\begin{array}{ll} \text{minimize} & b^T y \\ \text{subject to} & A^T y - z = c \\ & y, z \geq 0 \end{array}$$

Complementarity Conditions

$$\begin{array}{ll} x_j z_j = 0 & j = 1, 2, \dots, n \\ w_i y_i = 0 & i = 1, 2, \dots, m \end{array}$$

Matrix Notation

Can't write $xz = 0$. The product xz is undefined.

Instead, introduce a new notation:

$$x = \begin{bmatrix} x_1 \\ x_2 \\ \vdots \\ x_n \end{bmatrix} \implies X = \begin{bmatrix} x_1 & & & \\ & x_2 & & \\ & & \ddots & \\ & & & x_n \end{bmatrix}$$

Then the complementarity conditions can be written as:

$$\begin{aligned} XZe &= 0 \\ WYe &= 0 \end{aligned}$$

Optimality Conditions

$$\begin{aligned}Ax + w &= b \\ A^T y - z &= c \\ ZXe &= 0 \\ WYe &= 0 \\ w, x, y, z &\geq 0\end{aligned}$$

Ignore (temporarily) the nonnegativities.

$2n + 2m$ equations in $2n + 2m$ unknowns.

Solve'em.

Hold on. Not all equations are linear.

It is the nonlinearity of the complementarity conditions that makes LP fundamentally harder than solving systems of equations.

The Interior-Point Paradigm

Since we're ignoring nonnegativities, it's best to replace complementarity with μ -complementarity:

$$\begin{aligned}Ax + w &= b \\A^T y - z &= c \\ZXe &= \mu e \\WYe &= \mu e\end{aligned}$$

Start with an arbitrary (positive) initial guess: x, y, w, z .

Introduce *step directions*: $\Delta x, \Delta y, \Delta w, \Delta z$.

Write the above equations for $x + \Delta x, y + \Delta y, w + \Delta w$, and $z + \Delta z$:

$$\begin{aligned}A(x + \Delta x) + (w + \Delta w) &= b \\A^T(y + \Delta y) - (z + \Delta z) &= c \\(Z + \Delta Z)(X + \Delta X)e &= \mu e \\(W + \Delta W)(Y + \Delta Y)e &= \mu e\end{aligned}$$

Paradigm Continued

Rearrange with “delta” variables on left and drop nonlinear terms on left:

$$\begin{aligned}A\Delta x + \Delta w &= b - Ax - w \\A^T\Delta y - \Delta z &= c - A^T y + z \\Z\Delta x + X\Delta z &= \mu e - ZXe \\W\Delta y + Y\Delta w &= \mu e - WYe\end{aligned}$$

This is a **linear** system of $2m + 2n$ equations in $2m + 2n$ unknowns.

Solve'em.

Dampen the step lengths, if necessary, to maintain positivity.

Step to a new point:

$$\begin{aligned}x &\longleftarrow x + \theta\Delta x \\y &\longleftarrow y + \theta\Delta y \\w &\longleftarrow w + \theta\Delta w \\z &\longleftarrow z + \theta\Delta z\end{aligned}$$

(θ is the scalar damping factor).

Paradigm Continued

Pick a smaller value of μ for the next iteration.

Repeat from beginning until current solution satisfies, within a tolerance, optimality conditions:

primal feasibility $b - Ax - w = 0$.

dual feasibility $c - A^T y + z = 0$.

duality gap $b^T y - c^T x = 0$.

Theorem.

- Primal infeasibility gets smaller by a factor of $1 - \theta$ at every iteration.
- Dual infeasibility gets smaller by a factor of $1 - \theta$ at every iteration.
- If primal and dual are feasible, then duality gap decreases by a factor of $1 - \theta$ at every iteration (if $\mu = 0$, slightly slower convergence if $\mu > 0$).

LOQO

Hard/impossible to “do” an interior-point method by hand.

Yet, easy to program on a computer (solving large systems of equations is routine).

LOQO implements an interior-point method.

Setting option `loqo_options 'verbose=2'` in AMPL produces the following “typical” output:

LOQO Output

variables: non-neg 1350, free 0, bdd 0, total 1350
constraints: eq 146, ineq 0, ranged 0, total 146
nonzeros: A 5288, Q 0
nonzeros: L 7953, arith_ops 101444

| | Primal | | Dual | | Sig | |
|------|----------------|----------|---------------|----------|-----|--------|
| Iter | Obj Value | Infeas | Obj Value | Infeas | Fig | Status |
| 1 | -7.8000000e+03 | 1.55e+03 | 5.5076028e-01 | 4.02e+01 | | |
| 2 | 2.6725737e+05 | 7.84e+01 | 1.0917132e+00 | 1.65e+00 | | |
| 3 | 1.1880365e+05 | 3.92e+00 | 4.5697310e-01 | 2.02e-13 | | DF |
| 4 | 6.7391043e+03 | 2.22e-01 | 7.2846138e-01 | 1.94e-13 | | DF |
| 5 | 9.5202841e+02 | 3.12e-02 | 5.4810461e+00 | 1.13e-14 | | DF |
| 6 | 2.1095320e+02 | 6.03e-03 | 2.7582307e+01 | 4.15e-15 | | DF |
| 7 | 8.5669013e+01 | 1.36e-03 | 4.2343105e+01 | 2.48e-15 | | DF |
| 8 | 5.8494756e+01 | 3.42e-04 | 4.6750024e+01 | 2.73e-15 | 1 | DF |
| 9 | 5.1228667e+01 | 8.85e-05 | 4.7875326e+01 | 2.59e-15 | 1 | DF |
| 10 | 4.9466277e+01 | 2.55e-05 | 4.8617380e+01 | 2.86e-15 | 2 | DF |
| 11 | 4.8792989e+01 | 1.45e-06 | 4.8736603e+01 | 2.71e-15 | 3 | PF DF |
| 12 | 4.8752154e+01 | 7.26e-08 | 4.8749328e+01 | 3.36e-15 | 4 | PF DF |
| 13 | 4.8750108e+01 | 3.63e-09 | 4.8749966e+01 | 3.61e-15 | 6 | PF DF |
| 14 | 4.8750005e+01 | 1.81e-10 | 4.8749998e+01 | 2.91e-15 | 7 | PF DF |
| 15 | 4.8750000e+01 | 9.07e-12 | 4.8750000e+01 | 3.21e-15 | 8 | PF DF |

OPTIMAL SOLUTION FOUND

A Generalizable Framework

Start with an optimization problem—in this case LP:

$$\begin{array}{ll} \text{maximize} & c^T x \\ \text{subject to} & Ax \leq b \\ & x \geq 0 \end{array}$$

Use slack variables to make all inequality constraints into non-negativities:

$$\begin{array}{ll} \text{maximize} & c^T x \\ \text{subject to} & Ax + w = b \\ & x, w \geq 0 \end{array}$$

Replace nonnegativity constraints with *logarithmic barrier terms* in the objective:

$$\begin{array}{ll} \text{maximize} & c^T x + \mu \sum_j \log x_j + \mu \sum_i \log w_i \\ \text{subject to} & Ax + w = b \end{array}$$

Incorporate the equality constraints into the objective using *Lagrange multipliers*:

$$L(x, w, y) = c^T x + \mu \sum_j \log x_j + \mu \sum_i \log w_i + y^T (b - Ax - w)$$

Set derivatives to zero:

$$\begin{aligned} c + \mu X^{-1}e - A^T y &= 0 && \text{(deriv wrt } x) \\ \mu W^{-1}e - y &= 0 && \text{(deriv wrt } w) \\ b - Ax - w &= 0 && \text{(deriv wrt } y) \end{aligned}$$

Introduce *dual complementary variables*:


$$z = \mu X^{-1}e$$

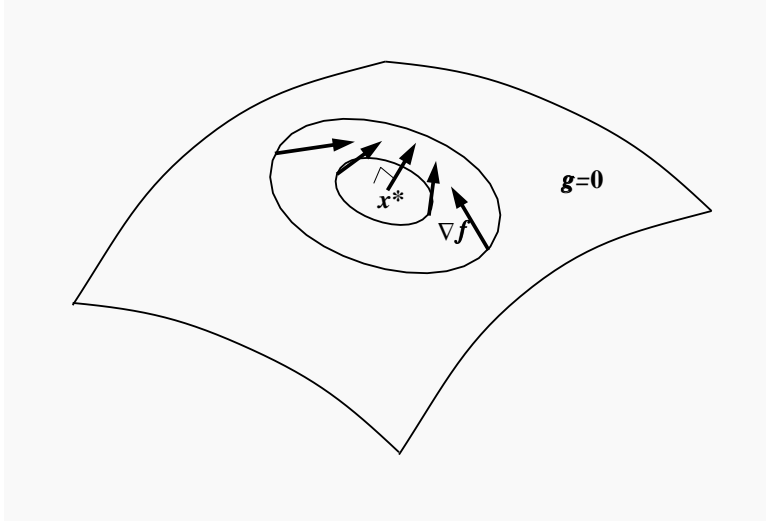
Rewrite system:

$$\begin{aligned} c + z - A^T y &= 0 \\ XZ e &= \mu e \\ WY e &= \mu e \\ b - Ax - w &= 0 \end{aligned}$$

Logarithmic Barrier Functions

Plots of $\mu \log x$ for various values of μ :





$$\begin{aligned} & \text{maximize} && f(x) \\ & \text{subject to} && g_1(x) = 0 \\ & && g_2(x) = 0 \end{aligned}$$

