About MITACS

Scientific Programs

International Program

Events & News

Events

Current News

Media Releases

Media Coverage

Reports

Connections - MITACS Newsletter

Student Notes - SAC Newsletter

Accolades

Media Kit

Subscribe

Photos

Commercialization

Partnerships

Member Resources

Contact Us
Subscribe
Search
Print
Home
En Français

Networks of Centres of Excellence
Réseaux de centres d'excellence

High Performance Optimization: Theory, Algorithm Design and Engineering Applications

by Dr. Tamás Terlaky (McMaster University) and Dr. Miguel Anjos (University of Waterloo)

Two New York Times front-page articles -the first on November 7, 1979, and the second on November 19, 1984 - heralded breakthroughs in the "solution of mathematical problems by computer" and a "breakthrough in problem solving" respectively. The "problems" referred to in both cases are linear optimization problems, and the fact that these two developments made the headlines strikingly illustrates the importance of linear optimization in many businesses.

A typical optimization problem consists of three parts. First, there are one or more decision variables representing the decisions to be made. Second, there is an objective to optimize (maximize profit or minimize cost); this objective is a function of the decision variables. Third, there are restrictions on the values that the decision variables can take; these restrictions are the problem's constraints. The choice of values for the decision variables that satisfies the constraints and gives the best possible value of the objective function is the optimal solution. A simple example of an optimization problem is the so-called diet problem, which asks for the cheapest combination of foods that will satisfy all the daily nutritional requirements of a person. In this optimization problem, the decision variables are the amounts of each food to be eaten, the objective is to minimize the total cost of the diet, and the constraints require that the food amounts be chosen so that the diet contains minimum amounts of nutrients such as calories, vitamins and minerals, as well as staying within prescribed maximum amounts of, say, fat and sodium.

Linear optimization (LO) refers to optimization problems with an objective given by a linear function and with constraints expressed by linear equalities and inequalities. The diet problem can be formulated as a LO problem, and was studied in the 1940s to help minimize the costs of the US Army. Nowadays, LO is an ubiquitous tool in many industries, including transportation, manufacturing, retail, and finance.

Linear optimization problems have the property that the set of possible solutions is always a polyhedron. A familiar example of a polyhedron is a cube, and it can be proven that if we wish to optimize any linear objective function over the cube (in any dimension), then the optimal solution will always be found by evaluating the objective function at all the vertices (corner points) of the cube, and choosing the vertex giving the best objective value. Until the 1980s, LO problems were mostly solved using the simplex method, which moves on the boundary of the polytope from vertex to vertex until the optimal solution is found. Unfortunately, the number of vertices of the cube is exponential in the dimension of the problem; for example, in three dimensions, there are 23 = 8 vertices. This means that the simplex method can take a very long time to find the optimal solution. In theory, this difficulty holds for general LO problems, but in practice, the simplex method only checks a relatively small number of vertices before finding the best one.

This fact has made it extremely popular and useful in practice.

The 1979 headline announced the proof by Leonid Khachiyan that the so-called ellipsoid method could solve LO problems in a polynomial number of steps, i.e., much more efficiently than the worst-case exponential number of steps of the simplex method. Unfortunately, the ellipsoid method was found to perform quite poorly in practice. But only five years later, the 1984 headline announced Narendra Karmarkar's results, and heralded the so-called interior-point revolution that has unfolded in the past two decades.

Unlike the simplex method, both the ellipsoid and interior-point methods follow a sequence of points inside the polyhedral set that ultimately leads to the optimal vertex. This revolution has led to major advancements not only for LO, but also for all areas of optimization. These advancements in turn led to significant advances in the applicability of optimization in many real-world contexts by allowing the rapid computation of solutions to problems that were previously considered impossible to solve.

The MITACS project High Performance Optimization: Theory, Algorithm Design and Engineering Applications is at the forefront of these advances. Indeed, its research is advancing both the theory and implementation of optimization algorithms as well as improving their applications in several areas of engineering. In this overview, we mention some of the project's contributions to the theory, algorithms and computational practice of optimization, as well as the contributions being made to two areas of engineering, namely Very Large Scale Integration (VLSI) and Electricity Markets.

A key ingredient of interior-point methods (IPMs) for LO is the barrier parameter that uniquely defines a path leading to the optimal solution set, and decreases to zero as the method converges to an optimal solution. Broadly speaking, there are two kinds of IPMs: small-update methods, which decrease the barrier parameter slowly, and large-update methods, which decrease it more aggressively. While small-update methods are theoretically faster than large-update methods, the latter converge more rapidly in practice. Many attempts to reconcile these facts had only limited success. However, progress in this regard has the potential to yield significant algorithmic improvements, and hence to have major impact for the users of these algorithms. In recent years, researchers from this project have investigated a new paradigm for IPMs. Two important consequences of this paradigm are: a significant narrowing of the gap between the theoretical and practical performance of IPMs, and several types of IPMs that are efficient both in theory and in practice.

The advent of IPMs opened vast new opportunities for a broad and powerful class of problems, namely Semidefinite Optimization (SDO) problems. SDO can be viewed as a generalization of LO that allows the decision variables to be matrices. Much research in the last decade has focused on analyzing and implementing efficient IPMs for SDO, so that several excellent solvers are now available. This has resulted in a vast increase in the size of SDO problems that can be solved, and thus expanded its practical applicability. The Advanced Optimization Laboratory (AdvOL) at McMaster University, where this MITACS project is based, is responsible for maintaining and developing SeDuMi, one of the most popular SDO solvers. The development of SeDuMi is now an essential part of this project. Since its inauguration in December 2004, the SeDuMi web page has attracted almost 250 000 visitors and supported thousands of downloads of the software.

We now turn our attention to some of the engineering areas in which this project has made significant contributions, beginning with Very Large Scale Integration (VLSI) chip design. Modern integrated circuits contain millions of transistors, and as circuits increase in size, complexity and functionality, the combinatorial optimization problems of partitioning the chip into modules, placing the modules, and routing the connections between the modules become increasingly complex. Traditional heuristics, such as simulated annealing and taboo-search, tend to become trapped in local minima and provide no means to measure how close the obtained solution is to optimality. On the other hand, methods that seek optimal solutions are disadvantaged by the exponential growth of both model size and solution time.

One thrust of research in this project is the use of SDO techniques to obtain the first non-trivial global lower bounds for some VLSI placement problems. Even if the optimal solution is not known, lower bounds provide a measure of the distance of any solution to optimality, and hence provide an estimate of the quality of solutions obtained by any algorithm. Another thrust is the development of Fast Approximation Algorithms (FAAs) for some of these problems. FAAs are algorithms that are guaranteed to be efficient, and to provide solutions provably close to optimality. For the routing problem in VLSI, researchers in this project made several theoretical advances using both SDO and FAAs. Both the SDO and FAA approaches have been implemented and tested computationally on well-known VLSI benchmark test sets. Preliminary results are extremely encouraging, and current research is developing more sophisticated implementations of these techniques, as well as extending them to more sophisticated VLSI design problems, such as the use of multi-layered designs.

A new addition to this project is the area of electricity markets. This is of great relevance to Canadian society today. Until recently, the generation, transmission and distribution of electricity in Canada was owned mostly by provincial governments, and the Canadian Hydros were responsible for the planning, construction and operation of the power systems that are the backbone of the electricity business. Significant deregulation in North America began in the 1990s, and led to the recent implementation of competitive electricity markets in parts of Canada, in particular Alberta (2001) and Ontario (2002). The main motivation behind this thrust has been the need to improve efficiencies and reduce costs through the competition generated by free markets.

The problems in some electricity markets in North America, including the recent major blackouts (e.g. the August 2003 blackout in North America's north-east), have brought to the forefront of public debate the strategic issue of securing reliable and affordable electricity supplies in Canada. In particular, one of the main reasons for the problems in most electricity markets is the lack of demand responsiveness to grid conditions and price variations which too often lead to serious power shortage and price volatility problems. The deregulation and privatization of electricity markets to make them more competitive has led to the need for new analysis and operating tools that properly dispatch supply and demand powers at optimal prices while guaranteeing adequate system security levels. Optimization is thus at the core of most of the technical and economic decision-making tools being used in energy markets. For example, system operators assign certain amounts of electricity to generators by solving a LO problem that minimizes the cost of meeting the demand for electricity, subject to constraints that include stability of the overall power network. Improved modeling techniques, together with the progress in optimization algorithms mentioned above, are beginning to have a profound effect in demonstrating how electricity markets can be managed with the greatest benefit to all parties involved, including Canadian consumers.

Finally, we want to mention that all the research being carried out in this project would not be possible without the constant advice and support, both financial and technical, of our industrial partners: Bell Canada (via Bell University Laboratories at the University of Waterloo), ABB Canada, Hydro One Networks, the Independent Energy System Operator (IESO) of Ontario, CITO, IBM, Rogers and the Princess Margaret Hospital in Toronto.

We hope that this overview provides the readers of Connections with a taste of the variety of applications of optimization, and of the breadth of research activities supported by MITACS through this project. At the present time, the area of optimization is truly an exciting area of research, and we look forward to many more years of exciting development in optimization and engineering, with perhaps even a newspaper headline!

previous article newsletter highlights **next** article