Business and Mathematics: A Saga of 25 Years of Progress in Optimization

Robert E. Bixby

GUROBI
OPTIMIZATION

A Short Bio

- 1972: Ph.D. Operations Research, Cornell
- 1972-1985: Very theoretical research in OR
- 1980: IBM PCs introduced, began computational work
  - 1983-84: Began work on CPLEX
- 1988: Founded CPLEX optimization
  - Cofounder: Janet Lowe
  - 1992: Market leader, displacing IBM
  - 1997: Sold to ILOG for $30M, value of deal $60M
- 2008: Founded Gurobi optimization
  - Cofounders: Zonghao Gu, Ed Rothberg

Outline

- Progress in Optimization
  - Short history of Linear Programming
    - Computational advances
  - Mixed-integer programming
    - What it is and why it's useful
    - How MIP is perceived in many communities
    - Computational progress
- A story of two companies
  - CPLEX Optimization
  - Gurobi Optimization
- Closing words of wisdom

Linear Programming

Linear Programming History

- George Dantzig, 1947
  - Introduced LP
  - Invented “primal” simplex algorithm.
  - First LP solved: Laderman (1947), 9 cons., 77 vars., 120 MAN-DAYS
  - Four Nobel Prizes in LP

- Commercial Developments
  - 1951: first computer implementations of simplex algorithm
  - 1960: LP becomes commercially viable (oil industry)
  - 1972: IBM MPSX/370 introduced, dominates commercial landscape
  - Mid 1980s: market stagnated, LP considered a “done subject”

They were WRONG: LP was far from finished

A Production Planning Model

401,640 cons. 1,584,000 vars. 9,498,000 nonzeros

Solution time line **(2.0 GHz P4)**:		**Speedup**
◦ 1988 (CPLEX 1.0):	29.8 days	**1x**
◦ 1997 (CPLEX 5.0):	1.5 hours	**480x**
◦ 2003 (CPLEX 9.0):	59.1 seconds	**43500x**

LP Today: LP is considered a solved problem by practitioners.

Mixed-Integer Programming

A Definition

A *mixed-integer program* (MIP) is an optimization problem of the form

$$\begin{aligned} \textit{Minimize} \quad & c^T x \\ \textit{Subject to} \quad & Ax = b \\ & l \leq x \leq u \\ \text{some or all } & x_j \text{ integer} \end{aligned}$$

Customer Applications (2013)

- Accounting
- Advertising
- Agriculture
- Airlines
- ATM provisioning
- Compilers
- Defense
- **Electrical power**
- **Energy**
- **Finance**
- Food service
- Forestry
- Gas distribution
- Government
- Internet applications
- **Logistics/supply chain**
- Medical
- Mining
- National research labs
- Online dating
- Portfolio management
- Railways
- Recycling
- Revenue management
- Semiconductor
- Shipping
- Social networking
- Sourcing
- Sports betting
- Sports scheduling
- Statistics
- Steel Manufacturing
- Telecommunications
- Transportation
- Utilities
- **Workforce Management**

Solving MIPs

The Perception

- Electrical Power
  - ERPI GS-6401, June 1989: Mixed-integer programming (MIP) is a powerful modeling tool, "They are, however, theoretically complicated and computationally cumbersome". In other words: MIP is an interesting "toy", but it just doesn't work in practice.
- Computer Science
  - "MIP is NP-hard. Really the only choice is to use heuristics."
- Story from INFORMS San Francisco
  - Several management science experts (as reported by a Gurobi user)
    - "MIP? Do people still work on that? I thought that was abandoned in the 80s?"

This Perception is Just Plain WRONG

Example: Supply-chain scheduling

A Consulting Engagement: Early 1997

- Model description:
  - Weekly model, daily buckets: Objective to minimize end-of-day inventory.
  - Production (single facility), inventory, shipping (trucks), wholesalers (demand known)
- Initial modeling phase
  - Simplified prototype + complicating constraints (production run grouping req't, min truck constraints)
  - RESULT: Couldn't get good feasible solutions.
- Decomposition approach
  - Talk to current scheduling team: They first decide on "producibles" schedule. Simulate using heuristics.
  - Fixed model: Fix variables and run MIP

Example: Supply-chain scheduling (cont.)

CPLEX 5.0 (1997):

```
Integer optimal solution (0.0001/0):  Objective =    1.5091900536e+05
Current MIP best bound =    1.5090391809e+05 (gap = 15.0873)
Solution time = 3465.73 sec.  Iterations = 7885711  Nodes = 489870 (2268)
```

CPLEX 11.0 (2007):

```
Implied bound cuts applied:  60
Flow cuts applied:  85
Mixed integer rounding cuts applied:  41
Gomory fractional cuts applied:  29

MIP - Integer optimal solution:  Objective =  1.5091900536e+05
Solution time =    0.63 sec.  Iterations = 2906  Nodes = 12
```

Original model: CPLEX 11 solved to optimality in 15 minutes (20% improvement in solution quality).

MIP Computational History: 1950 -1998

- 1954 Dantzig, Fulkerson, S. Johnson: 42 city TSP
  - Solved to optimality using LP and cutting planes
- 1957 Gomory
  - Cutting plane algorithms
- 1960 Land, Doig; 1965 Dakin
  - B&B
- 1964-68 LP/90/94
  - First commercial application
- IBM 360 computer
  - 1974 MPSX/370
  - 1976 Sciconic
    - LP-based B&B
    - MIP became commercially viable
- 1975 – 1998 Good B&B remained the state-of-the-art in commercial codes, in spite of ….
  - Edmonds, polyhedral combinatorics
  - 1973 Padberg, cutting planes
  - 1973 Chvátal, revisited Gomory
  - 1974 Balas, disjunctive programming
  - 1983 Crowder, Johnson, Padberg: PIPX, pure 0/1 MIP
  - 1987 Van Roy and Wolsey: MPSARX, mixed 0/1 MIP
  - TSP, Grötschel, Padberg, …

1998 … A New Generation of MIP Codes

- Linear programming
  - Stable, robust dual simplex
- Variable/node selection
  - Influenced by traveling salesman problem
- Primal heuristics
  - 12 different tried at root
  - Retried based upon success
- Node presolve
  - Fast, incremental bound strengthening (very similar to Constraint Programming)
- Presolve – numerous small ideas
  - Probing in constraints:

    $\sum x_j \leq (\sum u_j)\, y, \quad y = 0/1$

    ➔ $x_j \leq u_j y$ (for all j)
- Cutting planes
  - Gomory, mixed-integer rounding (MIR), knapsack covers, flow covers, cliques, GUB covers, implied bounds, zero-half cuts, path cuts

Some Test Results – from 2008

- Test set: 1852 real-world MIPs
  - Full library
    - 2791 MIPs
  - Removed:
    - 559 "Easy" MIPs
    - 348 "Duplicates"
    - 22 "Hard" LPs (0.8%)
- Parameter settings
  - Pure defaults
  - 30000 second time limit
- Versions Run
  - CPLEX 1.2 (1991) -- CPLEX 11.0 (2007)

MIP Speedups

CPLEX MIP Speedups 1991-2008

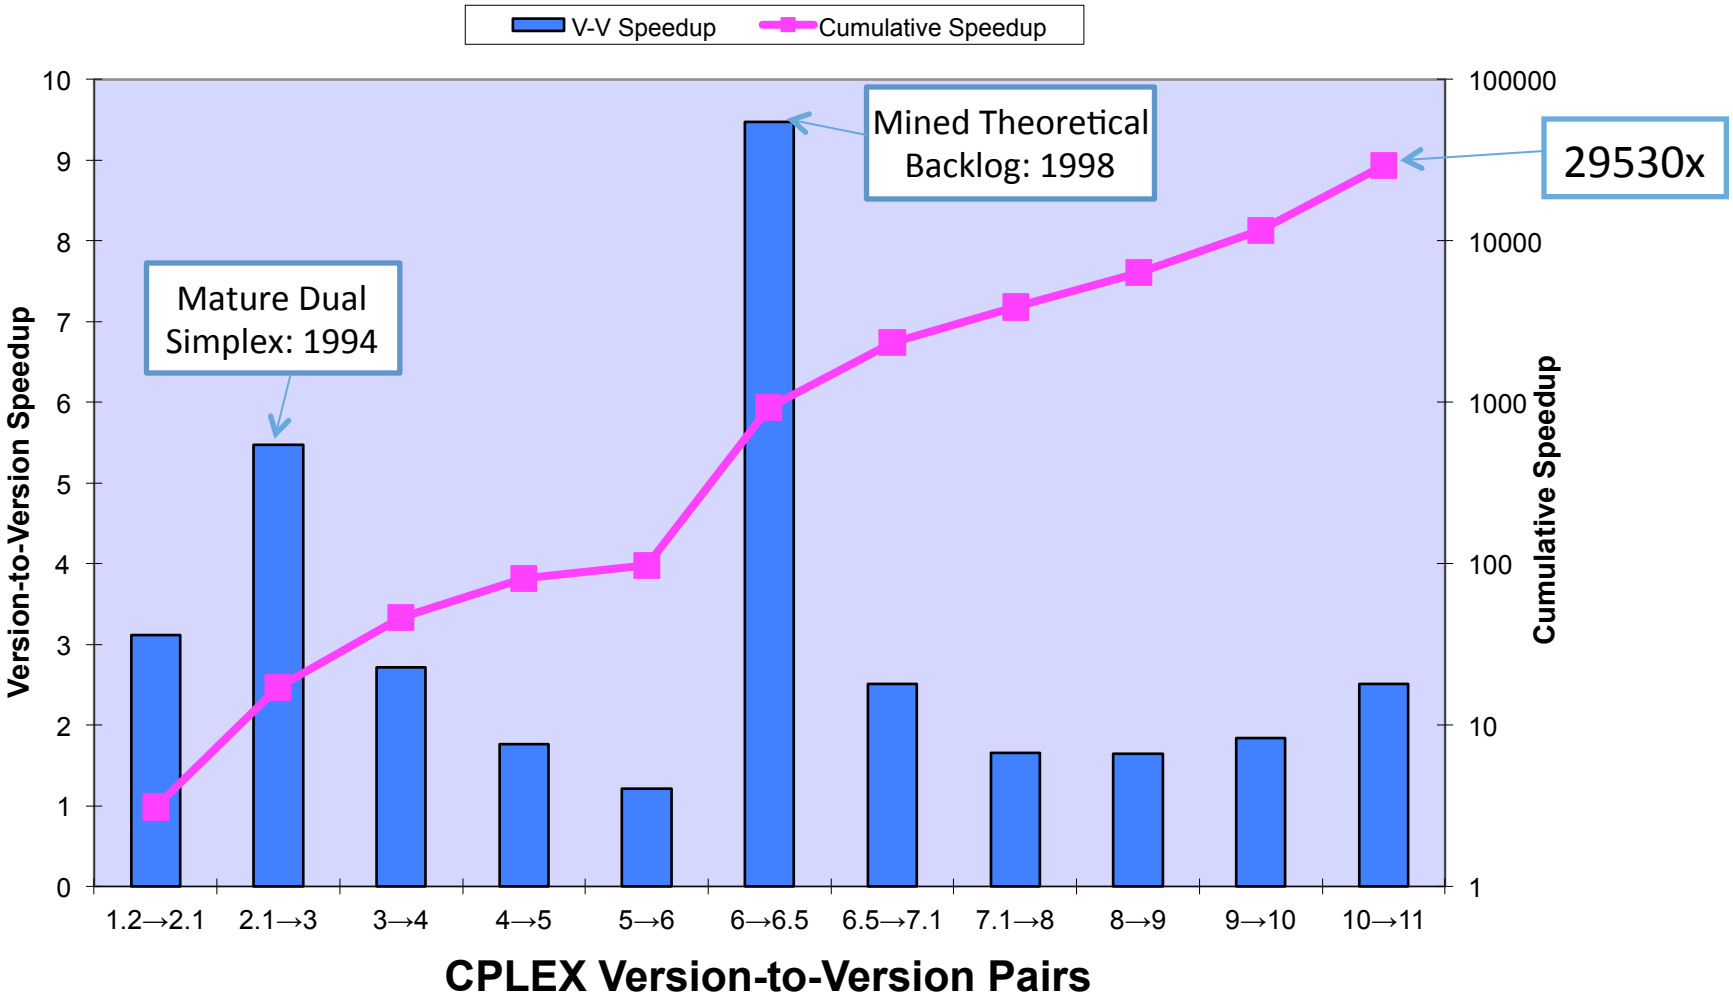

Gurobi MIP Speedups 2009-2014

- Gurobi 1.0 and CPLEX 11.0 roughly equivalent
- Gurobi version-to-version improvements
  - Gurobi 1.0 -> 6.0: 29.4X
- **Overall improvement:** 1990 to 2014
  - **Algorithms**: 870,000x
  - **Machines**: 6,500x
  - NET: Algorithm x Machine 5,600,000,000x
    (180 years / 5.6B ~= 1 second)

Founding Two Companies

CPLEX Optimization: The Start

- 1982-1987: Initial code development
  - Wrote classroom code
  - Code adopted by Chesapeake Decisions Sciences
    - Early supply chain company, founded by Tom Baker
    - Really remained a hobby
  - Purchased by Amaco 1986
  - *I realized I had something of value, and that I was pretty good at this odd mixture of mathematics, numerics, and computer science*

- 1988: CPLEX Optimization incorporated
  - Cofounded with Janet Lowe

CPLEX Optimization: Why was it successful?

- Willingness to take **RISKS**
  - I realized I didn't know how to start a business (1988)
  - I invited Janet and Todd Lowe to join, and had to give them ½ the business
- Developed a VISION
  - I realized that important work was going on, and this work needed an embeddable solver
  - This led to the notion of a *callable library and an object oriented design*, and this was a perfect fit for the industrial applications of the future
- Developed new commercial opportunities for solvers
  - The OEM model for solvers: embed your solver in other software, and let other companies do the selling for you
- Business Execution
  - We ran circles around IBM on customer care (support, licensing, pricing)
- Performance and *numerical stability*
  - CPLEX was just a good solver
- **Result:**
  - **By the late 1990s, CPLEX was completely dominant in its space**
  - **Sold to ILOG for $30M**

Gurobi Optimization: Back to Basics

- Cofounded is 2008
  - Our goal: Make our customers successful with optimization
- Back to Basics
  - Business execution
    - The best product support
    - The most flexible pricing and licensing
  - The best possible performance
    - Get the best people, and focus on this goal
  - Product innovation
    - Cloud offering, client-server offering, distributed computing
- Where we are now
  - Growing faster than CPLEX (revenues exceed those after 10 years)
  - 17 people (CPLEX had 13 when it was sold in 1997)
  - Distributors in four countries
  - Branch office in Germany
  - The top three developers in the industry

Gurobi: Three applications

- NFL
  - 2014 schedule was computed using Gurobi
  - They "solve" a single, very-difficult MIP
  - Computed 24x7 from the day after the Super Bowl until schedule release
    - Running on 800 cores
  - Produced 5-10 "good" schedules from which the owners chose one.

- New York ISO (NYISO)
  - Manages electrical-power generation in greater NY area (~1500 units)
  - Solve day-ahead unit-commitment, day-of-operation dispatch, and spot-pricing models
    - Running on 100 clients, 6 servers (using client-server model)
    - Pricing LP is solved every 5 minutes throughout the day

- eHarmony
  - Once per week eHarmony computes and "optimal" world-wide matching
  - The underlying MIP has 115 million variables and 40 million constraints
    - Solve time: approximately two hours
  - Model was built by a psychology PhD!!

Building a Successful Business

Three key ideas:

- Be prepared to take calculated risks
- Get the right people
- Objectively evaluate your markets