Mixed Integer Programming



Mixed Integer Programming

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

$$\begin{array}{ll} \underset{x}{\operatorname{minimize}} & \sum_{\substack{j=1\\n}}^{n} c_j x_j \\ \text{subject to} & \sum_{\substack{j=1\\\ell_j \leq x_j \leq u_j, \quad j=1,\ldots,n,}^{n} A_{ij} x_j = b_i, \quad i = 1,\ldots,m, \\ & \ell_j \leq x_j \leq u_j, \quad j = 1,\ldots,n, \\ & \text{some or all } x_j \text{ integer} \end{array}$$

Why do we care about this problem?



Applications

- 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

3

- National research labs
- Online dating
- Portfolio management
- Railways
- Recycling
- Revenue management
- Semiconductor
- Shipping
- Social networking
- Sports betting
- Sports scheduling
- Statistics
- Steel Manufacturing
- Telecommunications
- Transportation
- Utilities
- Workforce scheduling





National Football League



© 2015 Gurobi Optimization

NFL Schedule

5

- NFL revenues (2014) from TV contracts: ~\$5 billion in 2014
- > 256 games in regular season
- Each team plays 16 games (with 1 BYE) over 17 weeks
- Each team's opponents are predetermined before scheduling
- NFL begins scheduling immediately after Superbowl (once all opponents known)
- Input from all 32 clubs (stadium blocks, travel considerations, competitive factors, etc.)
- Input from all 5 television partners (key games, dates, markets, etc.)
- Entire process takes 2–3 months



2013 NFL Schedule

DAL	NYG	PHL	WAS	СНІ	DET	GB	MIN	ATL	CAR	NO	ТВ	ARZ	STL	SF	SEA		BUF	MIA	NE	NYJ	BAL	CIN	CLE	PIT	HOU	IND	JAX	TEN	DEN	КС	OAK	SD
NYG	DAL	WAS	PHI	CIN	MIN	<u>SF</u>	DET	NO	SEA	ATL	LLA	STL	ARZ	GB	CAR	1	NE	CLE	BUF	тв	DEN	СНІ	MIA	TEN	SD	OAK	КС	PIT	BAL	JAC	IND	HOU
КС	DEN	SD	GB	MIN	ARZ	WAS	СНІ	STL	BUF	тв	NO	DET	ATL	SEA	SF	2	CAR	IND	LLN	NE	CLE	PIT	BAL	CIN	TEN	MIA	ОАК	HOU	<u>NYG</u>	DAL	JAC	PHI
STL	CAR	кс	DET	РІТ	WAS	CIN	CLE	MIA	NYG	ARZ	NE	NO	DAL	IND	JAC	3	LLA	ATL	тв	BUF	HOU	GB	MIN	СНІ	BAL	<u>SF</u>	SEA	SD	ОАК	PHI	DEN	TEN
SD	кс	DEN	OAK	DET	CHI		PIT	NE		MIA	ARZ	тв	SF	STL	HOU	4	BAL	NO	ATL	TEN	BUF	CLE	CIN	MIN	SEA	JAC	IND	LLA	PHI	NYG	WAS	DAL
DEN	PHI	NYG		NO	GB	DET		LAN	ARZ	СНІ		CAR	JAC	HOU	IND	5	CLE	BAL	CIN	ATL	MIA	NE	BUF		SF	SEA	STL	кс	DAL	TEN	SD	ОАК
WAS	СНІ	тв	DAL	NYG	CLE	BAL	CAR		MIN	NE	PHI	SF	HOU	ARZ	TEN	6	CIN		NO	РІТ	GB	BUF	DET	LLA	STL	SD	DEN	SEA	JAC	OAK	кс	IND
PHI	MIN	DAL	CHI	WAS	CIN	CLE	NYG	тв	STL		ATL	SEA	CAR	TEN	ARZ	7	MIA	BUF	LLA	NE	<u>PIT</u>	DET	GB	<u>BAL</u>	КС	DEN	SD	SF	IND	HOU		JAC
DET	PHI	NYG	DEN		DAL	MIN	GB	ARZ	тв	BUF	CAR	ATL	SEA	JAC	STL	8	NO	NE	MIA	CIN		NYJ	кс	OAK			SF		WAS	CLE	PIT	
MIN		OAK	SD	GB		CHI	DAL	CAR	ATL	LLAN	SEA		TEN		тв	9	КС	CIN	еп	NO	CLE	MIA	BAL	NE	IND	HOU		STL		BUF	PHI	WAS
NO	OAK	GB	MIN	DET	CHI	PHI	WAS	SEA	SF	DAL	MIA	HOU	IND	CAR	ATL	10	PIT	ТВ			CIN	BAL		BUF	ARZ	STL	TEN	JAC	<u>SD</u>		NYG	DEN
	GB	WAS	PHI	BAL	PIT	NYG	SEA	тв	NE	<u>SF</u>	ATL	JAC		NO	MIN	11	LLA	SD	CAR	BUF	CHI	CLE	CIN	DET	OAK	TEN	ARZ	IND	кс	DEN	HOU	MIA
<u>NYG</u>	DAL		SF	STL	ТВ	MIN	GB	NO	MIA	ATL	DET	IND	CHI	WAS		12		CAR	DEN	BAL	LLA		PIT	CLE	JAC	ARZ	HOU	OAK	NE	SD	TEN	КС
OAK	WAS	ARZ	NYG	MIN	GB	DET	CHI	BUF	тв	SEA	CAR	PHI	SF	STL	NO	13	ATL	LLA	HOU	MIA	PIT	SD	JAC	BAL	NE	TEN	CLE	IND	КС	DEN	DAL	CIN
CHI	<u>SD</u>	DET	кс	DAL	PHI	ATL	BAL	GB	NO	CAR	BUF	STL	ARZ	<u>SEA</u>	<u>SF</u>	14	тв	PIT	CLE	OAK	MIN	IND	NE	MIA	JAC	CIN	HOU	DEN	TEN	WAS	LLA	NYG
GB	SEA	MIN	ATL	CLE	BAL	DAL	PHI	WAS	LLN	STL	SF	TEN	NO	тв	NYG	15	JAC	NE	MIA	CAR	DET	PIT	СНІ	CIN	IND	HOU	BUF	ARZ	SD	OAK	кс	DEN
WAS	DET	CHI	DAL	PHI	NYG	<u> PIT</u>	CIN	SF	NO	CAR	STL	SEA	ТВ	ATL	ARZ	16	MIA	BUF	BAL	CLE	NE	MIN	LLN	GB	DEN	КС	TEN	JAC	HOU	IND	SD	ОАК
PHI	WAS	DAL	NYG	GB	MIN	СНІ	DET	CAR	ATL	тв	NO	SF	SEA	ARZ	STL	17	NE	LAN	BUF	MIA	CIN	BAL	PIT	CLE	TEN	JAC	IND	HOU	ОАК	SD	DEN	кс

NBC, CBS, FOX, ESPN, NFLN

Home Away

6

Left NFC (E, N, S, W), Right AFC (E, N, S, W)

GUROBI

OPTIMIZATION

Partial list of rules and goals

- Teams playing in London play home week prior, or have BYE week after
- Bills game in Toronto Week 13 against ATL, BAL, CAR, KC, MIA or NYJ
- No team has earliest BYE in consecutive seasons
- No team plays more than two road games against team coming off their BYE
- All teams playing road Thursday games are home the previous week
- All teams playing home Thursday games have limited travel previous week
- Minimize non-division games during hurricane season (in order to "swap" sites)
- Minimize 3-game road trips (and 3game home stands for teams with ticket issues)

- Minimize number of division series that end in first half of season
- Minimize number of games that would conflict with MLB postseason
- Maximize the number of late-season division games
- Minimize early-season games 1PM games for teams with weather concerns
- Minimize number of Pacific time zone teams that play at 1PM ET
- CBS/FOX have at least 3 1PM games every week, preferably geographically diverse
- CBS/FOX have at least 5 total games
- Team can play no more than 6 prime time games, and only 3 teams per year can play 6. All other teams can play no more than 5 prime time games (max 4 NBC)



Stadium blocks

DAL	NYG	PHL	WAS	СНІ	DET	GB	MIN	ATL	CAR	NO	тв	ARZ	STL	SF	SEA		BUF	MIA	NE	NYJ	BAL	CIN	CLE	PIT	HOU	IND	JAX	TEN	DEN	кс	OAK	SD
BB		BB SUN													BB SU-MO	1					BB SU-MO	BB SU-MO								BB SUN	BB SUN	BB SUN
BB		BB		COLL.	BB								LEASE			2					BB			BB	LEASE					BB	BB	
BB		BB			BB		HOME	PGA					LEASE		BB	3					1110	BB	COLL.	HOME		NO THU				BB	BB	BB
BB	BON	SUN			TH-SU			TOUR							BB	4	<u> </u>			BON	BB	BB	FB						<u> </u>	SUN	TH-SU	BB
TH-SU	JOVI			<u> </u>			LONDON								SUN	4				JOVI	TH-SU	SUN		LONDON								THU*
DIV		DIV			DIV		BYE			NO THU						5					DIV	DIV		BYE	LEASE					NASCAR	DIV	DIV
LCS		LCS		Marathon	1				NASCAR				LEASE		AUTO	6		CARNIVAL				LCS		LCS							LCS	
LCS		LCS											NO THU	AWAY		7						LCS		LCS							LCS	
WS		WS											Marathon	LONDON	MLS	8						ws		ws			LONDON				ws	
										NO THU				BYE		9											BYE					
													LEASE		MLS?	10									LEASE							
			SOCCER							NO THU			LEASE			11		NO MNF								NO PRIME						
															MLS?	12										LEASE						NO MNF
TGIV					TGIV					NO SUN 1p						13								HIGH SCHOOL								
KU HOOPS								SEC	ACC	NO SUN 1p			LEASE		MLS CUP	14	Toronto?								LEASE	BIG 10						
		ARMY NAVY								NO SUN			LEASE?			15																
																16																
									BELK							17																

- Of the 544 potential sites, almost 20% unavailable or had some sort of conflict
- 13 teams had 4 or more potential conflicts, affecting 25% or more of their potential home dates



A \$5 billion MIP Model?

- All constraints captured in one MIP:
 - original: 25K rows, 20K cols, 800K nz
 - presolved: 6K rows, 6K cols, 160K nz
 - 3600 binary variables
- What does a solution look like:
 - 256 binary variables to set to 1
 - Each captures the particulars of one game
 - the rest set to 0





New York Independent System Operator







What is NYISO?

- NYISO is a non-profit, statutory agency
 - began operating in 1999
- NYISO primary responsibilities:
 - operating New York's high-voltage grid
 - administering and monitoring state's wholesale electricity markets
 - long-term planning for high-voltage grid
 - reliability, future maintenance and expansion
- Size
 - grid encompasses 11,000 miles of transmission lines
 - manages 500 power-generation units
 - administers \$7.5 billion in annual transactions





Optimization Plays a Central Role In NYISO Operations

- Day-ahead problem
 - unit-commitment for the next day: which units to be used
 - markets close 5 am
 - commitments must be posted by 11 am
 - 30 minutes available to compute commitments
- Real-time problem
 - solved every 5 minutes, restricted unit-commitment
- Real-time dispatch (EASY!)
 - linear program, solved every 5 minutes
 - determines generator clearing prices



Traditional Solution Approach: Lagrangian Relaxation (LR)

- In use since the early 1980s
- Claimed to produce high-quality solutions
- Solution times are quite good
- Considerable reluctance to make a change
- However in 1999 MIP was demonstrated to be a viable alternative*
 - pressure was mounting to switch

*Rutgers academic-industry meeting sponsored by Federal Energy Regulatory Commission (FERC).



Experience in Other ISOs*

- In 2004, PJM implemented MIP in its day-ahead market
 - estimated of annual production cost savings of \$60 million
- In 2006, PJM implemented MIP in its real-time market lookahead
 - test findings of \$100 million in annual savings
- In April 2009, CAISO implemented MIP as part of its Market Redesign and Technology Update
 - estimated savings of \$52 million
- In 2009, the Southwest Power Pool (SPP) introduced MIP enhancements to its day-ahead market
 - estimated **\$103 million** in annual benefits

*FERC Staff Report, 2011 (Recent ISO Software Enhancements and Future Software and Modeling Plans)





MIP: Other Advantages

- Maintainability
 - LR code: 100,000 lines of FORTRAN
 - understood by only one or two people
- Transparency
 - MIP formulation has much simpler representation
 - easy to read, interpret, and maintain
 - solutions much easier to defend against legal challenges
- Extensibility
 - extremely difficult to add constraints to LR
 - new constraints easy to add to MIP formulation
 - combined cycle-generator modeling enabled



Solution Implementation

- Solution launch date: Q2 2014
- Computing environment
 - 50 HPUX client machines
 - 12 Linux compute servers (3-production, 3-stage, 3-development)
- Estimated savings: \$4M yearly production costs
- Future
 - underlying models are highly nonlinear
 - MIP speed improvements can always be used to improve model accuracy
 - conclusion:
 - continued MIP improvements will lead directly to efficiency gains in electricalpower markets!



Conclusion: Use MIP Instead of Custom Solution!

- Easier to maintain a MIP model than complex source code
- Easier to extend a MIP model than to adapt a custom algorithm
- Easier to demonstrate correctness of a MIP model
- Benefit from performance improvements over time coming from MIP solver vendors
 - algorithmic speed-up exceeded hardware speed-up during the last 20 years!

Solving Mixed Integer Programs



© 2015 Gurobi Optimization

MIP Application Types

- Static MIP
 - Formulate problem
 - Solve it with a black-box MIP algorithm
 - Read solution
 - Potentially adjust problem and iterate
 - most frequent use of MIP in practical applications
- Branch-and-cut
 - Problem has too many constraints to formulate in static fashion
 - e.g., classical TSP model: exponentially many sub-tour elimination constraints
 - Construct partial problem
 - Add violated constraints on demand
 - "Lazy constraint" separation callback: cut off primal infeasible LP or IP solutions
- Branch-and-price
 - Problem has too many variables to formulate in static fashion
 - e.g., many public transport and airline problems are solved via B&P
 - Pricing callback: cut off dual infeasible LP solutions
 - Usually needs problem specific branching rule that is compatible with pricing
 - Heuristic variant: column generation
 - Only apply pricing for the root LP, then solve static MIP with resulting set of variables



MIP Building Blocks

- Presolve
 - Tighten formulation and reduce problem size
- Solve continuous relaxations
 - Ignoring integrality
 - Gives a bound on the optimal integral objective
- Cutting planes
 - Cut off relaxation solutions
- Branching variable selection
 - Crucial for limiting search tree size
- Primal heuristics
 - Find integer feasible solutions



MIP Building Blocks

Presolve

- Tighten formulation and reduce problem size
- Solve continuous relaxations
 - Ignoring integrality
 - Gives a bound on the optimal integral objective
- Cutting planes
 - Cut off relaxation solutions
- Branching variable selection
 - Crucial for limiting search tree size
- Primal heuristics
 - Find integer feasible solutions



LP Presolve

Goal

- Reduce the problem size
 - Speedup linear algebra during the solution process

Example

$x + y + z \le 5$	(1)
u - x - z = 0	(2)
$0 \le x, y, z \le 1$	(3)
u is free	(4)

- Reductions
 - Redundant constraint
 - (3) \Rightarrow x + y + z \leq 3, so (1) is redundant
 - Substitution
 - (2) and (4) \Rightarrow u can be substituted with x + z



MIP Presolve

- Goals:
 - Reduce problem size
 - Speed-up linear algebra during the solution process
 - Strengthen LP relaxation
 - Identify problem sub-structures
 - Cliques, implied bounds, networks, disconnected components, ...
- Similar to LP presolve, but more powerful:
 - Exploit integrality
 - Round fractional bounds and right hand sides
 - Lifting/coefficient strengthening
 - Probing
 - Does not need to preserve duality
 - · We only need to be able to uncrush a primal solution
 - Neither a dual solution nor a basis needs to be uncrushed





MIP Presolve

- Goals:
 - Reduce problem size
 - Speed-up linear algebra during the solution process
 - Strengthen LP relaxation
 - Identify problem sub-structures
 - Cliques, implied bounds, networks, disconnected components, ...
- Similar to LP presolve, but more powerful:
 - Exploit integrality
 - Round fractional bounds and right hand sides
 - Lifting/coefficient strengthening

model	W	ithout pre	esolve	with presolve							
model	rows	cols	LP obj	rows	cols	LP obj					
roll3000	2291	1166	11097.1	987	855	11120.0					
neos-787933	1897	236376	3.0	41	126	30.0					



Single-Row Reductions

- Clean-up rows
 - Discard empty rows
 - Discard redundant inequalities: $sup\{A_{r}, x\} \le b_r$
 - Remove coefficients with tiny impact $|a_{ij} \cdot (u_j l_j)|$
- Bound strengthening
 - $\ \ \, \text{ a}_{rj} > \text{ 0, s} := b_r \text{ inf}\{A_r.x\} \Longrightarrow x_j \leq I_j + s/a_{rj}$
 - $\label{eq:arg} \ \ \ a_{rj} < 0 \text{, } \text{s} \text{:= } b_r \text{ } \text{inf} \{A_r.x\} \Longrightarrow x_j \geq u_j + \text{s}/a_{rj}$
- Coefficient strengthening for inequalities
 - $\circ~j\in I,~a_{rj}>0,~t{:=}~b_r^{}-sup\{A_{r\cdot}^{}x\}+a_{rj}>0$

$$\Rightarrow a_{rj} := a_{rj} - t$$
, $b_r := b_r - u_j t$





Single-Row Reductions – Performance



OPTIMIZATION

Single-Row Reductions – Performance



benchmark data based on Gurobi 5.6





Single-Column Reductions

- Remove fixed variables and empty columns
 - If $|u_j l_j| \le \epsilon$, fix to some value in $[l_j, u_j]$ and move terms to rhs
 - Choice of value can be very tricky for numerical reasons
- Round bounds of integer variables
- Strengthen semi-continuous and semi-integer variables
- > Dual fixing, substitution, and bound strengthening
 - Variable x_i does not appear in equations

$$\label{eq:cj} \circ \ c_j \geq 0 \text{, } A_{\cdot j} \geq 0 \ \Rightarrow \ x_j \mathrel{\mathop:}= I_j$$

- $c_j \ge 0, A_{,j} \ge 0$ except for $a_{ij} < 0,$ $z = 0 \rightarrow row i redundant,$ $z = 1 \rightarrow x_j = u_j$ $\Rightarrow x_j := l_j + (u_j - l_j) \cdot z$
- $\circ \ c_j \geq 0 \text{, all rows i with } a_{ij} < 0 \text{ redundant for } x_j \geq t \ \Rightarrow \ x_j \leq max\{l_j,t\}$



Single-Column Reductions – Performance



benchmark data based on Gurobi 5.6

© 2015 Gurobi Optimization

GUROBI

OPTIMIZATION

Multi-Row Reductions

- Parallel rows
 - Search for pairs of rows such that coefficient vectors are parallel to each other
 - Discard the dominated row, or merge two inequalities into an equation
- Sparsify
 - Add equations to other rows in order to cancel non-zeros
 - Can also add inequalities with explicit slack variables
- Multi-row bound and coefficient strengthening
 - Like single-row version, but use other rows to get tighter bound on infimum and supremum \Rightarrow tighter bounds, better coefficients
- Clique merging
 - Merge multiple cliques into larger single clique, e.g.:



Multi-Row Reductions

31



benchmark data based on Gurobi 5.6



Multi-Column Reductions

- Fix redundant penalty variables
 - Penalty variables: support(A.j) = 1
 - Multiple penalty variables in a single constraint
 - Some can be fixed if others can accomplish all that is needed
- Parallel columns (say, columns 1 and 2): $A_{.1} = sA_{.2}$
 - $u_2 = \infty$, $c_1 \ge sc_2$, $2 \notin I$ or $(|s| = 1, \{1,2\} \subseteq I)$: $x_1 := I_1$
 - $I_2 = -\infty$, $c_1 \le sc_2$, $2 \notin I$ or $(|s| = 1, \{1,2\} \subseteq I)$: $x_1 := u_1$
 - $c_1 = sc_2$, 1,2 \notin I or (|s| = 1, {1,2} \subseteq I): $x_{1'} := x_1 + sx_2$
 - Detection algorithm: two level hashing plus sorting
- Dominated columns: $A_{.1} \ge sA_{.2}$, only inequalities
 - $u_2 = \infty$, $c_1 \ge sc_2$, $2 \notin I$ or $(|s| = 1, \{1,2\} \subseteq I)$: $x_1 := I_1$
 - Detection algorithm: essentially pair-wise comparison
 - Can be very expensive: needs work limit



Multi-Column Reductions – Performance



benchmark data based on Gurobi 5.6

GUROBI

OPTIMIZATION

Full Problem Reductions

- Symmetric variable substitution
 - Integer variables in same orbit can be aggregated if the involved symmetries do not overlap
 - Continuous variables in same orbit can always be aggregated
 - Issue: symmetry detection can sometimes be time consuming!
- Probing
 - $\,\circ\,$ Tentatively fix binary x=0 and x=1 , propagate fixing to get domain reductions for other variables
 - $\begin{array}{lll} \cdot & x = 0 \rightarrow y \leq u_0, \, x = 1 \rightarrow y \leq u_1 & \Rightarrow \, y \leq max\{u_0, u_1\} & (bound strength.) \\ \cdot & x = 0 \rightarrow y = l_y, \, x = 1 \rightarrow y = u_y & \Rightarrow \, y := l_y + (u_y l_y) \cdot x & (substitution) \\ \cdot & ay \leq b, \, x = 1 \rightarrow ay \leq d < b & \Rightarrow \, ay + (b-d) \cdot x \leq b & (lifting) \end{array}$
 - Sequence dependent
 - Can be very time consuming
 - Needs specialized data structures and algorithms
- Implied Integer Detection
 - Primal version: ax + y = b, x integer variables, $a \in \mathbb{Z}^n$, $b \in \mathbb{Z} \Rightarrow y$ integer
 - Dual version:
 - One of the inequalities for y will be tight, but do not know which
 - If all those inequalities lead to primal version of implied integer detection, y is implied integer



Full Problem Reductions



benchmark data based on Gurobi 5.6



MIP Building Blocks

- Presolve
 - Tighten formulation and reduce problem size
- Solve continuous relaxations
 - Ignoring integrality
 - Gives a bound on the optimal integral objective
- Cutting planes
 - Cut off relaxation solutions
- Branching variable selection
 - Crucial for limiting search tree size
- Primal heuristics
 - Find integer feasible solutions


Primal and Dual LP

Primal Linear Program:
T

min $c^T x$

$$s.t. \quad Ax = b$$
$$x \ge 0$$

- Weighted combination of constraints (y) and bounds (z) yields $y^{T}Ax + z^{T}x \ge y^{T}b$ (with $z \ge 0$)
- Dual Linear Program:

$$\max \qquad y^{T}b \\ s.t. \qquad y^{T}A + z^{T} = c^{T} \\ z \geq 0$$

Strong Duality Theorem: $c^T x^* = y^{*^T} b$ (if primal and dual are both feasible)



Karush-Kuhn-Tucker Conditions

Conditions for LP optimality:

- Primal feasibility: Ax = b $(x \ge 0)$
- Dual feasibility: $A^{T}y + z = c$ ($z \ge 0$)
- Complementarity: $x^{T}z = 0$

	<u>Primal feas</u>	<u>Dual feas</u>	Complementarity
Primal simplex	Maintain	Goal	Maintain
Dual simplex	Goal	Maintain	Maintain
Barrier	Goal	Goal	Goal





Phase 1: find some feasible vertex solution





- Pricing: find directions in which objective improves and select one of them
 - Gurobi parameters: SimplexPricing, NormAdjust





Ratio test: follow outgoing ray until next vertex is reached





Iterate until no more improving direction is found



MIP – LP Relaxation





MIP – LP Relaxation



No feasible solutions can be better than an LP optimum



Side Note – Performance Variability

- Measuring performance of MIP is a difficult task
- Seemingly performance neutral changes can have a dramatic impact on the solve time for an individual model
 - Changing the random seed
 - Permuting the columns and rows
 - e.g., store model in *.lp file format and reading it back in
 - Solving model on a different machine
 - Different operating system
 - Different CPU
- Main reason: degeneracy of LP relaxation
 - Most models have multiple optimal solutions to their LP relaxation
 - LP solver's choice of solution is arbitrary
 - Depends on random seed, column and row ordering, tiny differences in numerics
 - Different optimal solution vector leads to different cutting planes, different branching decisions, and thus different search tree
- Conclusion: need big test set to measure performance
 - Use multiple random seeds to artificially increase test set
- More on performance variability: tomorrow



MIP Building Blocks

- Presolve
 - Tighten formulation and reduce problem size
- Solve continuous relaxations
 - Ignoring integrality
 - Gives a bound on the optimal integral objective
- Cutting planes
 - Cut off relaxation solutions
- Branching variable selection
 - Crucial for limiting search tree size
- Primal heuristics
 - Find integer feasible solutions



















© 2015 Gurobi Optimization







No feasible solutions can be better than an LP optimum



Cutting Planes – Overview

- General-purpose cutting planes
 - Gomory mixed integer cuts
 - Mixed Integer Rounding (MIR) cuts
 - Flow cover cuts
 - Lift-and-project (L&P) cuts
 - Zero-half and mod-k cuts
 - •
- Structural cuts
 - Implied bound cuts
 - Knapsack cover cuts
 - GUB cover cuts
 - Clique cuts
 - Multi-commodity-flow (MCF) cuts
 - Flow path cuts
 - ••••



Chvatal-Gomory Cuts

Consider a rational polyhedron

 $P=\{x\in R^n\mid Ax\leq b,\,x\geq 0\},\,A\in \mathbb{Q}^{mxn},\,b\in \mathbb{Q}^m$

- We want to find the integer hull $P_1 = \text{conv} \{ x \in \mathbb{Z}^n \mid Ax \le b, x \ge 0 \}$
- Chvatal–Gomory procedure:
 - 1. Choose non–negative multipliers $\lambda \in \mathbb{R}^{m}_{\geq 0}$
 - 2. Aggregated inequality $\lambda^T A x \leq \lambda^T b$ is valid for P because $\lambda \geq 0$
 - 3. Relaxed inequality $\lfloor \lambda^T A \rfloor x \leq \lambda^T b$ is still valid for P because $x \geq 0$
 - 4. Rounded Inequality $\lfloor \lambda^T A \rfloor x \leq \lfloor \lambda^T b \rfloor$ is still valid for P_I because $x \in \mathbb{Z}^n$
- CG procedure suffices to generate all non-dominated valid inequalities for P₁ in a finite number of iterations!
 - $P^{(0)} = P$
 - $P^{(k)} = P^{(k-1)} \cap \{CG \text{ cuts for } P^{(k-1)}\}: k-th CG \text{ closure of } P \text{ is a polyhedron!} \}$
 - CG rank of a valid inequality for P_I : minimum k s.t. inequality is valid for $P^{(k)}$

Gomory Fractional Cuts

- How to find the multipliers λ in the Chvatal–Gomory procedure?
- Read them from an optimal simplex tableau!

 $x_i + \lfloor (A_B^{-1})_i A_N \rfloor x_N \leq \lfloor (A_B^{-1})_i b \rfloor$

- Note: new slack variable is integer
 - Hence, procedure can be iterated
- Gomory (1963): These cuts lead to an integral LP solution in a finite number of iterations.
- Conclusion:
 - Integer programs with rational data can be solved by pure cutting plane algorithms in finite time
 - What about mixed integer programs?
 - And is this result useful in practice?



Gomory Cuts – Numerics



> Pictures stolen from Zanette, Fischetti and Balas (2011):

"Lexicography and degeneracy: can a pure cutting plane algorithm work?"

- Evolution of optimal basis matrix condition number and average absolute coefficient values of Gomory cuts for model "sentoy"
 - Gomory fractional cuts added in rounds

Mixed Integer Rounding Cuts

• Consider $S := \{(x,y) \in \mathbb{Z} \times \mathbb{R}_{\geq 0} \mid x - y \leq b\}.$

Then,
$$x - \frac{1}{1 - f_0} y \le \lfloor b \rfloor$$

is valid for S with $f_0 := b - \lfloor b \rfloor$.

- Example: $x y \le 2.5$
- MIR cut: $x 2y \le 2$





Mixed Integer Rounding Cuts

• Consider S := {
$$(x,y) \in \mathbb{Z} \times \mathbb{R}_{\geq 0} \mid x - y \leq b$$
}.

Then,
$$x - \frac{1}{1 - f_0} y \le \lfloor b \rfloor$$

is valid for S with $f_0 := b - \lfloor b \rfloor$.

• Consider S := {(x,y) $\in \mathbb{Z}_{\geq 0}^{p} \times \mathbb{R}_{\geq 0}^{q} \mid ax + dy \leq b$ }.

Then,

$$\sum \left(\lfloor a_i \rfloor + \frac{\max\{f_i - f_0, 0\}}{1 - f_0} \right) x_i + \sum \left(\frac{\min\{d_j, 0\}}{1 - f_0} \right) y_j \le \lfloor b \rfloor$$

is valid for S with $f_i := a_i - \lfloor a_i \rfloor$, $f_0 := b - \lfloor b \rfloor$.



Mixed Integer Rounding Cuts

- General idea just like Chvatal-Gomory cuts:
 - 1. Choose non–negative multipliers $\lambda \in \mathbb{R}^{m}_{_{\geq 0}}$
 - 2. Aggregated inequality $\lambda^T A x \leq \lambda^T b$ is valid for P because $\lambda \geq 0$
 - 3. Apply MIR formula to aggregated inequality to produce cutting plane
- Cut separation procedure of Marchand and Wolsey (1998, 2001):
 - 1. Start with one constraint of the problem (do this for each one), call this the "current aggregated inequality"
 - 2. Apply MIR procedure to current aggregated inequality
 - (a) Complement variables if LP solution is closer to upper bound
 - (b) For each a_i in constraint and each of $\delta \in \{1,2,4,8\}$ divide the constraint by $\delta |a_i|$
 - (c) Apply MIR formula to resulting scaled constraint
 - (d) Choose most violated cut from this set of MIR cuts
 - (e) Check if complementing one more (or one less) variable yields larger violation
 - 3. If no violated cut was found (and did not yet reach aggregation limit):
 - (a) Add another problem constraint to the current aggregated inequality such that a continuous variable with LP value not at a bound is canceled
 - (b) Go to 2



Gomory Mixed Integer Cuts

- Just an alternative way to aggregate constraints
- Read them from an optimal simplex tableau:
 - Let i be a basis index with $x_i^{\,\star} \not\in \mathbb{Z}$
 - Choose $\lambda^{T} = (A_{B}^{-1})_{i.}$
 - $^{\circ}$ Resulting aggregated inequality: $x_i^{} + (A_B^{-1})_{i\cdot}A_N^{} x_N^{} \leq (A_B^{-1})_{i\cdot}b$
- Apply MIR formula on resulting aggregated inequality
- In theory, always produces a violated cutting plane
- Practical issues:
 - Gomory Mixed Integer Cuts can be pretty dense
 - Numerics (in particular for higher rank cuts) can be very challenging
- But:
 - If done right, GMICs (together with MIRs) are currently the most important cutting planes in practice



Knapsack Cover Cuts

- A (binary) knapsack is a constraint $ax \le b$ with
 - $a_i \ge 0$ the weight of item i, i = 1, ..., n
 - $b \ge 0$ the capacity of the knapsack
- An index set $C \subseteq \{1,...,n\}$ is called a *cover*, if $\sum_{i=1}^{n} a_i > b$
- A cover C entails a cover inequality

$$\sum_{i \in C} x_i \le |C| - 1$$

Interesting for cuts: minimal covers

$$\sum_{i \in C} a_i > b \quad \text{and} \quad \sum_{i \in C'} a_i \le b \text{ for all } C' \subsetneq C$$



Knapsack Cover Cuts – Example

- Consider knapsack $3x_1 + 5x_2 + 8x_3 + 10x_4 + 17x_5 \le 24$, $x \in \{0,1\}^5$
- A minimal cover is $C = \{1, 2, 3, 4\}$
- Resulting cover inequality: $x_1 + x_2 + x_3 + x_4 \le 3$

- Lifting
 - If $x_5 = 1$, then $x_1 + x_2 + x_3 + x_4 \le 1$
 - $\circ \ \ \, \text{Hence, } x_1^{} \, + \, x_2^{} \, + \, x_3^{} \, + \, x_4^{} \, + \, 2 x_5^{} \, \leq \, 3 \, \, \text{is valid}$
 - Need to solve knapsack problem $\alpha_j := d_0 max\{dx \mid ax \le b a_j\}$ to find lifting coefficient for variable x_j
 - Use dynamic programming to solve knapsack problem



Cutting Planes – Performance



Achterberg and Wunderling: "Mixed Integer Programming: Analyzing 12 Years of Progress" (2013) benchmark data based on CPLEX 12.5

© 2015 Gurobi Optimization

MIP Building Blocks

- Presolve
 - Tighten formulation and reduce problem size
- Solve continuous relaxations
 - Ignoring integrality
 - Gives a bound on the optimal integral objective
- Cutting planes
 - Cut off relaxation solutions
- Branching variable selection
 - Crucial for limiting search tree size
- Primal heuristics
 - Find integer feasible solutions















© 2015 Gurobi Optimization









© 2015 Gurobi Optimization





© 2015 Gurobi Optimization







72




Solving a MIP Model

73





Branching Variable Selection

- Given a relaxation solution x*
 - Branching candidates:
 - Integer variables x_i that take fractional values
 - $x_j = 3.7$ produces two child nodes (x ≤ 3 or x ≥ 4)
 - Need to pick a variable to branch on
 - Choice is crucial in determining the size of the overall search tree



Branching Variable Selection

- What's a good branching variable?
 - Superb: fractional variable infeasible in both branch directions
 - Great: infeasible in one direction
 - Good: both directions move the objective
- Expensive to predict which branches lead to infeasibility or big objective moves
 - Strong branching
 - Truncated LP solve for every possible branch at every node
 - Rarely cost effective
 - Need a quick estimate



Pseudo-Costs

- Use historical data to predict impact of a branch:
 - Record $cost(x_i) = \Delta obj / \Delta x_i$ for each branch
 - Store results in a pseudo-cost table
 - Two entries per integer variable
 - Average down cost
 - Average up cost





Pseudo-Costs

- Use historical data to predict impact of a branch:
 - Record $cost(x_j) = \Delta obj / \Delta x_j$ for each branch
 - Store results in a pseudo-cost table
 - Two entries per integer variable
 - Average down cost
 - Average up cost





Pseudo-Costs

- Use historical data to predict impact of a branch:
 - Record $cost(x_j) = \Delta obj / \Delta x_j$ for each branch
 - Store results in a pseudo-cost table
 - Two entries per integer variable
 - Average down cost
 - Average up cost





Pseudo-Costs Initialization

- What do you do when there is no history?
 - E.g., at the root node
- Initialize pseudo-costs [Linderoth & Savelsbergh, 1999]
 - Always compute up/down cost (using strong branching) for new fractional variables
 - Initialize pseudo-costs for every fractional variable at root
- Reliability branching [Achterberg, Koch & Martin, 2005]
 - Do not rely on historical data until pseudo-cost for a variable has been recomputed r times



Branching Rules – Performance

80



benchmark data based on CPLEX 12.5

Achterberg, Koch, and Martin: "Branching Rules Revisited" (2005)

GUROBI

OPTIMIZATION

MIP Building Blocks

- Presolve
 - Tighten formulation and reduce problem size
- Solve continuous relaxations
 - Ignoring integrality
 - Gives a bound on the optimal integral objective
- Cutting planes
 - Cut off relaxation solutions
- Branching variable selection
 - Crucial for limiting search tree size
- Primal heuristics
 - Find integer feasible solutions



Primal Heuristics Explained on Twitter





OPTIMIZATION

Primal Heuristics

- Try to find good integer feasible solutions quickly
 - Better pruning during search due to better bound
 - Reach desired gap faster
 - Often important in practice: quality of solution after fixed amount of time
- Start heuristics
 - Try to find integer feasible solution, usually "close" to LP solution
 - Rounding heuristics: round LP solution to integral values
 - Potentially, try to fix constraint infeasibilities
 - Fix-and-dive heuristics: fix variables, propagate, resolve LP
 - Feasibility pump: push LP solution towards integrality by modifying objective
 - RENS: Solve sub-MIP in neighborhood of LP solution
- Improvement heuristics
 - Given integer feasible solution, try to find better solution
 - 1-Opt and 2-Opt: Modify one or two variables to get better objective
 - Local Branching: Solve sub-MIP in neighborhood of MIP solution
 - Mutation: Solve sub-MIP in neighborhood of MIP solution
 - Crossover: Solve sub–MIP in neighborhood of 2 or more MIP solutions
 - RINS: Solve sub-MIP in neighborhood of LP and MIP solution

Primal Heuristics – Performance



Berthold: "Primal Heuristics for Mixed Integer Programs" (2006) benchmark data based on SCIP 0.82b

© 2015 Gurobi Optimization



Primal Heuristics – Measuring Performance

- Is time to optimality a good measure to assess impact of heuristics?
 - Goal of heuristics is to provide good solutions quickly
 - Faster progress in dual bound due to additional pruning is only secondary
 - Often important for practitioners:
 - Find any feasible solution quickly to validate that model is reasonable
 - Find good solution in reasonable time frame

• Primal gap:
$$\gamma^p(\tilde{x}) = \frac{|c^T x^* - c^T \tilde{x}|}{max\{|c^T x^*|, |c^T \tilde{x}|\}}$$

• Primal gap function: $p(t) = \begin{cases} 1, \text{ if no incumbent until time } t \\ \gamma^p(\tilde{x}(t)), \text{ with } \tilde{x}(t) \text{ being incumbent at time } t \end{cases}$

• Primal integral:
$$P(T) = \int_{t=0}^{T} p(t)dt$$

Primal Integral







86

Primal Heuristics – Performance



Berthold (2014): "Heuristic algorithms in global MINLP solvers" benchmark data based on SCIP 3.0.2



Putting It All Together



© 2015 Gurobi Optimization





















© 2015 Gurobi Optimization













© 2015 Gurobi Optimization

OPTIMIZATION







GUROBI

OPTIMIZATION







Presolvir	Presolved: 987 rows, 855 columns, 19346 nonzeros Variable types: 211 continuous, 644 integer (545 binary)									
	Root relaxation: objective 1.112003e+04, 1063 iterations, 0.03 seconds									
	Noc Expl U	les Inexpl	Curr Obj E	rent Nod Depth In	e tInf	Objec Incumbent	ctive Bounds c BestBd	 Gap	Wa It/Noo	ork de Time
	0	0	11120.027	79 0	154	_	11120.0279	_	_	0s
	0	0	11526.891	.8 0	207	-	11526.8918	-	_	0s
	0	0	11896.971	0 0	190	-	11896.9710	-	-	0s
Cutting P	 H 0 0 1066 1097 1135 3416 5485	0 2 702 724 710 887 634	12448.768 12956.267 12671.828 12732.560 12839.988 12885.365	34 0 76 31 85 8 91 32 80 46 52 52	181 192 147 126 136 143	5890.000000 15890.0000 13087.0000 12890.0000 12890.0000 12890.0000	12448.7684 12448.7684 12629.5426 12671.8285 12727.1362 12780.7059 12829.0134	21.7% 21.7% 3.50% 3.17% 1.26% 0.85% 0.47%	- 37.2 41.6 44.6 49.7 54.5	0s 0s 5s 10s 15s 20s 25s
Branching										



Performance Impact of MIP Solver Components (CPLEX 12.5 or SCIP)



© 2015 Gurobi Optimization

OPTIMIZATION

Parallelization

- Parallelization opportunities
 - Parallel probing during presolve
 - Almost no improvement
 - Use barrier or concurrent LP for initial LP relaxation solve
 - Only helps for large models
 - Run heuristics or other potentially useful algorithms in parallel to the root cutting plane loop
 - Moderate performance improvements: 20-25%
 - Does not scale beyond a few threads
 - Solve branch-and-bound nodes in parallel
 - Main speed-up for parallel MIP
 - Performance improvement depends a lot on shape of search tree
 - Typically scales relatively well up to 8 to 16 threads



Parallelization

Parallelization issues

- Determinism
- Load balancing
- CPU heat and memory bandwidth
 - Additional threads slow down main thread
- Root node does not parallelize well
 - Sequential runtime of root node imposes limits on parallelization speed-up
 - Amdahl's law
- A dive in the search tree cannot be parallelized
 - Parallelization only helps if significant number of dives necessary to solve model

More on parallel MIP: tomorrow



Parallel MIP – Performance



Achterberg and Wunderling: "Mixed Integer Programming: Analyzing 12 Years of Progress" (2013) benchmark data based on CPLEX 12.5, models with ≥ 100 seconds solve time



105



