


Making Molehills out of Mountains:

A Guided Tour of Large-scale Discrete Optimization



Ted Ralphs

Department of Industrial and Systems Engineering
Lehigh University, Bethlehem, PA

<http://www.lehigh.edu/~tkr2>

Outline of Talk

- Introduction
- Basic Methodology
- Advanced Techniques
- Computation
- Current Outlook

Introduction

Mathematical Programming Models

- What does *mathematical programming* mean?
- Programming here means “planning.”
- Literally, these are “mathematical models for planning.”
- Also called *optimization models*.
- Essential elements
 - Decision variables
 - Constraints
 - Objective Function
 - Parameters and Data

Forming a Mathematical Programming Model

- The general form of a *mathematical programming model* is:

$$\begin{array}{l} \min f(x) \\ \text{s.t. } g_i(x) \left\{ \begin{array}{l} \leq \\ = \\ \geq \end{array} \right\} b_i \\ x \in X \end{array}$$

$X \subseteq \mathbb{R}^n$ is an (implicitly defined) set that may be discrete.

- A *mathematical programming problem* is a problem that can be expressed using a mathematical programming model (called the *formulation*).
- A single mathematical programming problem can be represented using many different formulations (**important!**).

Types of Mathematical Programming Models

- The type of mathematical programming model is determined mainly by
 - The form of the objective and the constraints.
 - The form of the set X .
- In this talk, we consider **linear models**.
 - The objective function is linear.
 - The constraints are linear.
 - Linear models are specified by cost vector $c \in \mathbb{R}^n$, constraint matrix $A \in \mathbb{R}^{m \times n}$, and right-hand side vector $b \in \mathbb{R}^m$ and have the form

$$\begin{aligned} \min \quad & c^T x \\ \text{s.t.} \quad & Ax \geq b \\ & x \in X \end{aligned}$$

Linear Models


- Generally speaking, linear models are easier to solve than more general types of models.
- If $X = \mathbb{R}^n$, the model is called a *linear program* (LP).
- Linear programming models can be solved effectively.
- If some of the variables in the model are required to take on integer values, the model is called a *mixed integer linear programs* (MILPs).
- MILPs can be extremely difficult to solve in practice.

Modeling with Integer Variables


- Why do we need **integer variables**?
- If a variable represents the quantity of a physical resource that only comes in **discrete units**, then it must be assigned an integer value.
- We can use **0-1 (binary) variables** for a variety of other purposes.
 - Modeling yes/no decisions.
 - Enforcing disjunctions.
 - Enforcing logical conditions.
 - Modeling fixed costs.
 - Modeling piecewise linear functions.
- The simplest form of ILP is a **combinatorial optimization problem** (COP), where all variables are binary.

Combinatorial Optimization


- A *combinatorial optimization problem* $CP = (E, \mathcal{F})$ consists of
 - A *ground set* E ,
 - A set $\mathcal{F} \subseteq 2^E$ of *feasible solutions*, and
 - A *cost function* $c \in \mathbb{Z}^E$ (optional).
- The *cost* of $S \in \mathcal{F}$ is $c(S) = \sum_{e \in S} c_e$.
- A *subproblem* is defined by $\mathcal{S} \subseteq \mathcal{F}$.
- Problem: Find a least cost member of \mathcal{F} .



Cost 1100



Cost 1105



Cost 1107

Example: Perfect Matching Problem

- We are given a set N of n people that need to be paired in teams of two.
- Let c_{ij} represent the “cost” of the team formed by persons i and j .
- We wish to minimize the overall cost of the pairings.
- The nodes represent the people and the edges represent pairings.
- We have $x_{ij} = 1$ if i and j are matched, $x_{ij} = 0$ otherwise.
- To simplify the presentation, we assume that $x_{ij} = 0$ if $i \geq j$.

$$\begin{aligned} \min \quad & \sum_{\{i,j\} \in N \times N} c_{ij} x_{ij} \\ \text{s.t.} \quad & \sum_{j \in N} x_{ij} = 1, \quad \forall i \in N \\ & x_{ij} \in \{0, 1\}, \quad \forall \{i, j\} \in N \times N, i < j. \end{aligned}$$

Example: Facility Location Problem

- We are given n potential **facility locations** and m customers that must be serviced from those locations.
- There is a fixed cost c_j of opening facility j .
- There is a cost d_{ij} associated with serving customer i from facility j .
- We have two sets of binary variables.
 - y_j is 1 if facility j is opened, 0 otherwise.
 - x_{ij} is 1 if customer i is served by facility j , 0 otherwise.


$$\begin{aligned} \min \quad & \sum_{j=1}^n c_j y_j + \sum_{i=1}^m \sum_{j=1}^n d_{ij} x_{ij} \\ \text{s.t.} \quad & \sum_{j=1}^n x_{ij} = 1 && \forall i \\ & \sum_{i=1}^m x_{ij} \leq m y_j && \forall j \\ & x_{ij}, y_j \in \{0, 1\} && \forall i, j \end{aligned}$$

Basic Methodology

How Hard Are These Problems?


- In practice, solving these problems to optimality can be **extremely difficult**.
- Consider the well-known **Traveling Salesman Problem**.
 - We are given a set of cities and a cost associated with traveling between each pair of cities.
 - We want to find the least cost route traveling through every city and ending up back at the starting city.
- The number of possible solutions for the **TSP** is $n!$ (that's **HUGE**).
- Can we reasonably use exhaustive enumeration?

Example Instance of the TSP





Optimal Solutions to the 48 City Problem



Primal Algorithms

- Many optimization algorithms use an **improving search paradigm**.
 - Find a feasible starting solution.
 - In each iteration, determine an *improving feasible direction* and move in that direction to get to a better solution.
 - The step size is determined by performing a line search.
- Can this be made to work for integer programming?
- What are the difficulties?


Implicit Enumeration

- *Implicit enumeration* techniques try to enumerate the solution space in an intelligent way.
- The most common algorithm of this type is *branch and bound*.
- Suppose F is the set of feasible solutions for some MILP and we wish to solve $\min_{x \in F} c^T x$.
- Consider a *partition* of F into subsets F_1, \dots, F_k . Then

$$\min_{x \in F} c^T x = \min_{1 \leq i \leq k} \{ \min_{x \in F_i} c^T x \}$$

- *Idea*: If we can't solve the original problem directly, we might be able to solve the smaller *subproblems* recursively.
- Dividing the original problem into subproblems is called *branching*.
- Taken to the extreme, this scheme is equivalent to complete enumeration.
- We avoid complete enumeration primarily by deriving *bounds* on the value of an optimal solution to each subproblem.

The Geometry of Integer Programming



Bounding

- A *relaxation* of an ILP is an auxiliary mathematical program for which
 - the feasible region contains the feasible region for the original ILP, and
 - the objective function value of each solution to the original ILP is not increased.
- Types of Relaxations
 - **Continuous relaxations**
 - * Most common continuous relaxation is the *LP relaxation*.
 - * Obtained by dropping some or all of the integrality constraints.
 - * Easy to solve.
 - * Initial bounds weak, but can be strengthened with valid inequalities.
 - * Other relaxations are possible using **semi-definite programming**, for instance.
 - **Combinatorial relaxations**
 - * Obtained by dropping some of the linear constraints.
 - * Violation of these constraints can then be penalized in the objective function (*Lagrangian relaxation*)
 - * Bound strength depends on what constraints are dropped.


Branch and Bound Algorithm

- We maintain a queue of *active* subproblems initially containing just the *root subproblem*.
- We choose a subproblem from the queue and solve a relaxation of it to obtain a *bound*.
- The result is one of the following:
 1. The relaxation is infeasible \Rightarrow *subproblem is infeasible*.
 2. We obtain a feasible solution for the *MILP* \Rightarrow subproblem solved (new upper bound??).
 3. We obtain an optimal solution to the relaxation that is not feasible for the *MILP* \Rightarrow *lower bound*.
- In the first two cases, we are *finished*.
- In the third case, we compare the lower bound to the global upper bound.
 - If it exceeds the upper bound, we discard the subproblem.
 - If not, we *branch* and add the resulting subproblems to the queue.


Branching

Branching involves partitioning the feasible region with hyperplanes such that:

- All optimal solutions are in one of the members of the partition.
- The solution to the current relaxation is not in any of the members of the partition.



Branch and Bound Tree



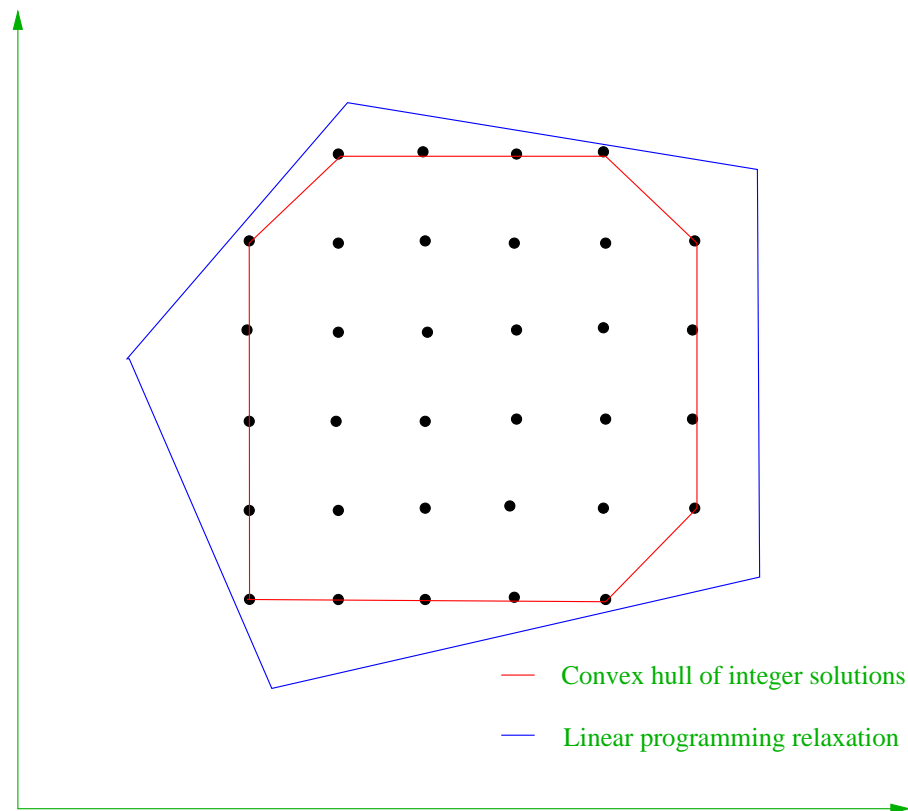
Advanced Techniques

Reformulation


- Not all formulations are created equal.
- A given mathematical programming problem may have many alternative formulations.
- For MILPs, “stronger” formulations make problems easier to solve.
- Example: Facility Location revisited
 - Consider the constraints $\sum_{i=1}^m x_{ij} \leq my_j \quad \forall j$.
 - These can be replaced with $x_{ij} \leq y_j \quad \forall i, j$
- Adding variables or recasting with a completely different set of variables can also help.
- Various *automatic reformulation techniques* have been successful in improving our ability to solve difficult problems.
- Decomposition, preprocessing, and generation of valid inequalities are three simple methods for strengthening initial formulations.

Generation of Valid Inequalities

Any inequality valid for all optimal solutions to a given MILP can be used to obtain improved bounds.



Valid Inequalities Examples



- Consider the graph on the left above.
- The **optimal perfect matching** has value $L + 2$.
- The optimal solution to the LP relaxation has value 3 .
- This formulation can be extremely **weak**.
- Add the **valid inequality** $x_{24} + x_{35} \geq 1$.
- Every perfect matching satisfies this inequality.

Decomposition Methods

- *Decomposition methods* try to reduce solution of a difficult model to solution of a series of easier models.
- This is done by relaxing certain constraints and then trying to implicitly enforce them.
- One typical approach, called *Lagrangian relaxation*, assigns a penalty in the objective function for violating the relaxed constraints.
- The difficult part is identifying which constraints to relax.
- Methods for automatically detecting structure and applying a decomposition method are a promising area for research.
- Methods for solving families of related integer programs are crucial to efficient implementations.

Preprocessing and Logical Tightening

- *Preprocessing* techniques use various logical rules to tighten bounds on both variables and constraints.
- Example: We can derive *implied bounds* for variables from each constraint $ax \leq b$. If $a_0 > 0$, then

$$x_1 \leq (b - \sum_{j:a_j>0} a_j l_j - \sum_{j:a_j<0} a_j u_j) / a_0$$

- Many other such rules can be applied in order to strengthen the formulation and obtain better bounds.
- Similar techniques are used in *constraint logic programming*.

Computation

Software Development

- How do practitioners use these tools?
 - As a “black box” to solve a given model.
 - As a “black box” embedded within a larger application using a callable library.
 - As building blocks within customized solvers.
- User interfaces
 - Modeling languages
 - Data interchange formats
 - Callable library
 - Software frameworks
- Software is the bridge between theory and practice.
- User interfaces allow practitioners to apply methodology developed by academics to *real problems*.

Large-scale Computation

- With large-scale computation, many practical issues arise.
 - speed
 - memory usage
 - numerical stability
- Dealing with these issues can be more of an art than a science.
- This is an area in which we have a great deal to learn.

Distributed Computing

- **High-performance computing** is becoming increasingly affordable.
- The use of parallel algorithms for solving large-scale problems has become a realistic option for many users.
- Developing parallel algorithms raises a range of additional issues.
- The name of the game in parallel computation is to **avoid doing unnecessary computations** (right hand doesn't know what the left hand is doing).
- To avoid unnecessary work, processing units have to **share information**.
- Information sharing also has a cost, so there is a tradeoff.
- Achieving the correct balance is challenging.
- This is an area of active research.

Applications


Application of the Traveling Salesman Problem

- Drilling Circuit Boards
- Continuous Line Drawing
- DNA Sequencing

DNA Sequencing and the TSP

- The *DNA sequencing problem* is to find the sequence of base pairs in a large fragment of DNA (length N).
- It is not practical to simply examine the DNA and determine the sequence.
- One approach is *sequencing by hybridization*: mix the DNA sequence with short oligonucleotides that bind with subsequences of the DNA.
- This results in an approximate list of all subsequences of length l that occur in the larger sequence.
- To reconstruct the original sequence, we simply have to correctly *order the subsequences*.
- This is easy if there are no errors.
- In the presence of errors, it becomes an optimization problem that is a variant of the TSP.

Current State of the Art



The Vehicle Routing Problem

In the **VRP**, we have additional constraints.

- There is a designated **depot** node (0).
- d is a vector of the customer **demands**.
- $V^- = V \setminus \{0\}$.
- k is the number of **routes**.
- C is the **capacity** of a truck.

A **feasible solution** is composed of:

- a **partition** $\{R_1, \dots, R_k\}$ of V such that $\sum_{j \in R_i} d_j \leq C$, $1 \leq i \leq k$;
- a **permutation** σ_i of $R_i \cup \{0\}$ specifying the order of the customers on route i .

ILP Formulation for the VRP

ILP Formulation:

$$\begin{aligned}\sum_{j=1}^n x_{0j} &= 2k \\ \sum_{j=1}^n x_{ij} &= 2 \quad \forall i \in V^- \\ \sum_{\substack{i \in S \\ j \notin S}} x_{ij} &\geq 2b(S) \quad \forall S \subset V^-, |S| > 1.\end{aligned}$$


$b(S)$ = lower bound on the number of trucks required to service S (nominally $\lceil (\sum_{i \in S} d_i) / C \rceil$).


If $C = \infty$, then we have the Multiple Traveling Salesman Problem.


How Hard is the VRP?

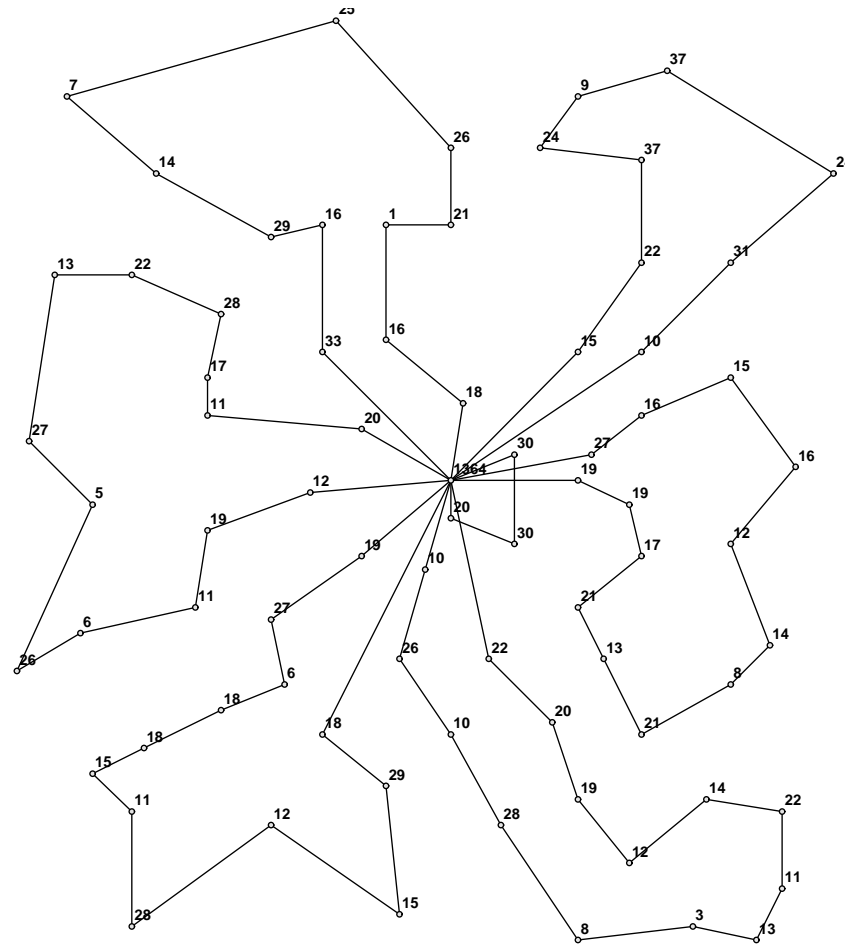
- Test Set
 - TSPLIB/VRPLIB
 - Augerat's repository
 - Available at BranchAndCut.org/VRP
- Largest VRP instance solved: F-n135-k7
- Largest TSP instance solved: d15112
- Smallest VRP instance unsolved (as of '01): B-n50-k8
- Time to solve B-n50-k8 as an MTSP: .1 sec
- Question: Why the gap?
- Answer: It is the intersection of two difficult problems.
 - Traveling Salesman Problem (Routing)
 - Bin Packing Problem (Packing)











The Set Partitioning Problem

- We are given a finite set S and subsets S_1, \dots, S_n of S , each with an associated cost $c(S_i), i \in [1..n]$.
- The *Set Partitioning Problem* is to determine $I \subset [1..n]$ such that $\bigcup_{i \in I} S_i = S$ and $\sum_{i \in I} c(S_i)$ is minimized.
- We can formulate this as an integer program.
 - Construct a $0 - 1$ matrix A in which $a_{ij} = 1$ if and only if the i^{th} element of S belongs to S_j .
 - The integer program is then

$$\begin{aligned} \min \quad & c^T x \\ \text{s.t.} \quad & Ax = 1 \\ & x \in \{0, 1\}^n \end{aligned}$$

- These integer programs are very difficult to solve.

The Set Partitioning Problem and Combinatorial Auctions

- A *combinatorial auction* is an auction in which participants are allowed to bid on subsets of the available goods.
- This accounts for the fact that some items have a greater worth when combined with other items.
- Example: FCC bandwidth auction
 - The FCC wishes to sell the licenses by bandwidth and region.
 - The value of a set of bandwidth licenses is increased if they are in contiguous regions.
- A set of items along with a price constitutes a *bid*.
- Given a set of bids, determining the *winners* of the auction is a set partitioning problem.

The Set Partitioning Problem and Vehicle Routing

- Problem data
 - A set of **customers** with known demands for a single commodity.
 - A fleet of identical **trucks** with fixed capacity.
 - A single fixed **depot**.
 - Travel times between pairs of locations.
- The *Vehicle Routing Problem* (VRP) is to
 - assign customers to trucks such that capacity is not exceeded and
 - sequence the customers assigned to each trucksuch that the total travel time is minimized.
- This can be formulated as a **set partitioning problem**.
 - The set to be partitioned is the **customers**.
 - The subsets are sets of customers that can be serviced by a single truck.
 - The cost for each subset is the cost of an optimal routing for that set of customers.

Resources (Shameless Self-Promotion)

- We have recently founded a laboratory devoted to computational optimization research called **COR@L** (coral.ie.lehigh.edu).
- Selected **COR@L** Software projects
 - SYMPHONY (www.branchandcut.org)
 - ALPS/BiCePS/BLIS
 - DECOMP
 - MINTO
 - MW
 - SUTIL
- The **COIN-OR Foundation** is a non-profit foundation founded by researchers from both industry and academia (www.coin-or.org).
 - We are dedicated to promoting the development and use of interoperable, open-source software for operations research.
 - We are also dedicated to **defining standards and interfaces** that allow software components to interoperate with other software and users.

Current Outlook

- Currently, IP researchers are in search of the “next big breakthrough” in methodology.
- More work is needed on improving user interfaces and making IP technology accessible to practitioners.
- There is still a lot to be learned about computation and large-scale IP.
- Applying IP to a new application area can have a big impact.
- Many application areas remain untapped.

My Research

- **Theory and Methodology**
 - Branch, cut, and price algorithms for large-scale MILP
 - Decomposition-based algorithms
 - Parallel algorithms
 - Multi-objective MILP
 - Sensitivity analysis and warm starting for MILP
- **Software Development**
 - **COIN-OR** (Open source software for operations research)
 - **SYMPHONY** (C library for parallel branch, cut, and price)
 - **ALPS/BiCePS/BLIS** (C++ library for scalable parallel search)
 - **DECOMP** (Framework for decomposition-based algorithms)
 - **VRP** Large-scale Vehicle Routing
 - **OSI** Open Solver Interface
- **Applications**
 - Logistics (routing and packing problems)
 - Network design
 - Combinatorial auctions