Modern Techniques in Discrete Optimization: Mathematics, Algorithms and Applications

Jesús A. De Loera (Univ. of California, Davis)
Jon Lee (Univ. of Michigan, Ann Arbor)

February 15, 2016

1 Overview of the Field

Discrete optimization is a vibrant area of the mathematical sciences devoted to finding optimal solutions given mathematical constraints that describe a finite or countable set of possible answers. For example, a classical problem in discrete optimization is the traveling salesman problem: For given n cities and the costs of traveling from city i to city j, we seek to find the cheapest route that visits each city once and returns to the starting city. Discrete optimization problems naturally arise in many kinds of applications including bioinformatics, telecommunications network design, airline and production scheduling, VLSI design, and efficient resource allocation, to name just a few. At the same time, the calculation of solutions requires sophisticated mathematics, for instance, there are deep connections between semidefinite optimization models and real algebraic geometry. The purpose of the workshop was to bring together researchers in the theory of discrete optimization who work in North America (only two researchers came from Spain). The goal was to strengthen the community of North American discrete-optimization researchers and to support the emerging group on this topic in Mexico.

2 Highlights of the Meeting

Three general research themes emerged in the discussions, we present only a few of the new mathematical advances we discussed during the meeting.

2.1 Mathematical Foundations

Combinatorial and algebraic structures are crucial in the modeling and solution of discrete optimization problems. In particular objects such as graphs and matroids appear in multiple contexts. A number of the presentations dealt with questions of fundamental structures. E.g., in the presentation of Dr. Araujo (UNAM, Mexico) there was a discussion of open problems about *cages*. A graph is k-regular if all the vertices have degree equal to k. The *girth* of a graph, denoted by g, is the number of edges in the smallest cycle(s). The Cage Problem is related to the existence and construction of regular graphs with a fixed girth and the minimum number of vertices possible. A $(k;g)$-graph is a k-regular graph of girth g, and a $(k;g)$-cage is a $(k;g)$-graph of minimum order. Cages were introduced by Tutte in 1947 and their existence was proved by Erdös and Sachs in 1963 for all integer values of k and g with $k \geq 2$. Since then most work carried out has focused on constructing cages with the smallest number of vertices. This graphs are important as the exemplify extreme symmetry and can serve as tests for algorithms.

Araujo and several other researchers have worked in this problem since 2006. She has constructed, jointly with several coauthors, different families of "small" $(k;g)$-graphs for many values of k and g. During the presentation, she explained the main techniques for constructing these graphs; using finite geometries. She also discussed the minimal values for the upper bounds for the order of the cages known up to date. Of great importance for this conference Araujo presented specific open problems, several participants made suggestions on how to use computers to carry on construction of $(k;5)$-minimal cages.

Speaker Criel Merino (UNAM, Mexico) discussed some new work about matroids (matroids are abstractions of the nation of vector space generated by a matrix and of the notion of graph). Traditionally, the focus from a graph theory perspective about the chromatic polynomial has been to find the positive integer roots λ, which correspond to the graph not being properly colourable with λ colours. A growing body of work has begun to emerge in recent years more concerned with the behavior of real or complex roots of the chromatic polynomial. Perhaps one of the outstanding open questions concerning real zeros is to determine tight bounds on the largest real zero of the chromatic polynomial. One such bound was given by A. Sokal, and more recently by F. M. Dong, and depends on the maximum vertex degree of the graph. The corresponding invariant in Matroids is the *characteristic polynomial*. In his presentation Merico proved that, for any prime power q and constant k, the characteristic polynomial of any loopless, $GF(q)$-representable matroid with tree-width k has no real zero greater than q^{k-1}. This is a remarkable result.

Many of the presentations and discussions had to do with the geometric fundamentals necessary to solve optimization problems. For example, *semi-definite* programming is a hot area of research, but to make progress one needs to understand the geometric structure of the cone $\mathcal{C}$ of positive semi-definite matrices. Speaker Greg Blekherman (Georgia Tech, USA), discussed subcones of low-rank inside the cone $\mathcal{C}$. More precisely a spectrahedral cone C is a slice of the cone of positive semidefinite matrices with a linear subspace L. The ranks of extreme rays of spectrahedral cones have been a subject of extensive study. It is natural to ask for what subspaces L do all of the extreme rays of C have rank 1? When L is a union of coordinate subspaces the answer was given by Agler-Helton-McCullough-Rodman. It turns out that this question has an unexpected connection to algebraic geometry and Blekherman presented a full classification of such spectrahedral cones based on the classification of small reduced schemes by Eisenbud-Green-Hulek-Popescu. It is a deep connection between two seemingly distinct areas of mathematics.

Not only geometry can be useful within discrete optimization, sometimes the benefit goes in the other direction too. One can approach difficult combinatorial questions using optimization. Thomas Rothvoss (Univ. of Washington, Seattle, USA), discussed his new work in the theory of discrepancy. A classical theorem of Spencer shows that any set system with n sets and n elements admits a coloring of discrepancy $O(n^1/2)$. Recent exciting work of Bansal, Lovett and Meka shows that such colorings can be found in polynomial time. In fact, the Lovett-Meka algorithm finds a half integral point in any "large enough" polytope. However, their algorithm crucially relies on the facet structure and does not apply to general convex sets. We show that for any symmetric convex set K with measure at least exp(-n/500), the following algorithm finds a point y in $K \cap [-1,1]^n$ with Omega(n) coordinates in $\{-1,+1\}$: (1) take a random Gaussian vector x; (2) compute the point y in $K \cap [-1,1]^n$ that is closest to x. (3) return y. This provides another truly constructive proof of Spencer's theorem and the first constructive proof of a Theorem of Gluskin and Giannopoulos.

Currently a very active subject of research is to understand the combinatorial geometry of polyhedra, in particular the combinatorial diameter of polyhedra. The *combinatorial diameter* of a polyhedron is the maximum number of edges (or 1-faces) needed to connect any two of its vertices. Alternatively, it can be defined as the diameter of the skeleton (or 1-skeleton) of the polyhedron. Motivated by the study of the worst-case performance of the Simplex algorithm to solve linear optimization problems, one considers the purely geometric problem of what is the largest possible combinatorial diameter of convex polytopes with given number of facets and dimension. One of the most famous statements associated with the combinatorial diameter is the *Hirsch conjecture*, stated in 1957 by Warren M. Hirsch. It claimed an upper bound of $f-d$ on the combinatorial diameter any d-dimensional polyhedron with f facets. It was finally disproved by Santos in 2010. But today Finding a good bound on the maximal diameter $D(d, f)$ of the vertex-edge graph of a polytope in terms of its dimension d and the number of its facets f is one of the basic open questions in geometry. There were a number of discussions and presentations around this subject. E.g., Antoine Deza from (McMaster University, Canada) Talked about the diameters of Lattice polytopes (that is those polytopes with integer coordinates in their vertices). He presented both older and recent results dealing with the diameter of lattice polytopes. In a similar spirit, Tamon Stephen (Simon Fraser Univ. Canada) Spoke about abstractions of the simplex method. In addition to the local edge directions used in the simplex method, a discrete set of additional directions are available. Some of these pass into the interior of the feasible region or of mid-dimensional faces. He and his collaborators studied the case where the available directions are the *circuits* or *elementary vectors*, i.e. the support minimal solutions to the homogenization of the equations defining the feasible region. These can also be thought of as potential edge directions of the system under varying right-hand sides of the equations. Once a direction is chosen, it is followed as far as feasibility allows.

A necessary condition for the existence of high-quality pivoting algorithms of this type is the existence of short paths between the vertices of the feasible region in the sense of using only a small number of these pivots and augmentations. This leads to the notion of the *circuit diameter* of a polyhedron P, which is number of pivots in the longest minimal path between any pair of vertices in P. The circuit diameter is an analogue of the *combinatorial diameter* of a polytope, which is the longest minimal path using only edge moves. Thus the circuit diameter is a lower bound for the combinatorial diameter, and the bound is tight for key families. Borgwardt, Finhold and Hemmecke (2014) investigated circuit diameter and showed that dual transportation polyhedra have a very low circuit diameter, lower than their combinatorial diameter. They asked if the Hirsch bound of $f-d$ (the number of facets of the polytope minus its dimension) which was conjectured as an upper bound on the combinatorial diameter, might hold for the combinatorial diameter. Stephen et al showed that some known non-Hirsch polyhedra, notably the Klee-Walkup polyhedron, are not counter-examples to this circuit Hirsch conjecture.

Also related to the graph of convex polyhedra (feasible regions of linear optimization problems) Walter Morris (George Mason Univ. USA) spoke about the connectivity property of graphs of *oriented matroids*. Oriented matroids are generalizations of convex polyhedra. Holt and Klee proved that if P is a d-dimensional polytope and g is a linear function on P that is not constant on any edge of P, there are d independent monotone paths from the source to the sink of the digraph defined by the vertices and edges of P directed according to the directions of increase of g. Mihalisin and Klee later proved that every orientation of the graph of a 3-polytope that is acyclic and admits 3 independent monotone paths from the source to the sink is obtained from some 3-polytope P and some linear function on P. But this is not true in general! Morris proved analogs of Mihalisin and Klee's theorem and the 3 and 4-dimensional versions of Holt and Klee's theorem for oriented matroid programs. Here acyclicity is replaced by the requirement that there be no directed cycle contained in a face of the polytope.

2.2 Models, reformulations and computational methods

The practical solution of discrete optimization problems depends on the use of structural mathematical properties to develop efficient algorithms. Often a problem is unsolvable directly, but one has to developed strategies to approximate it. Several talks in the meeting discussed topics of the algorithmic solution of various types of optimization problems.

Amitabh Basu (Johns Hopkins University, MD, USA) presented a survey of the latest advances to generate cuts in mixed integer programming problems (an old principle of work is that by adding linear conditions to the original formulation one can remove fractional solutions). Basu gave a rather useful introduction to *cut-generating functions*. Cut-generating functions are means to have “a priori” formulas for generating cutting planes for general mixed-integer optimization problems. Let S be a closed subset of $\mathbb{R}^n$ with $0 \notin S$. Consider the following set, parametrized by matrices R, P:

$$X_S(R,P) := \left\{(s,y) \in \mathbb{R}^k_+ \times \mathbb{Z}^\ell_+ \ : \ Rs + Py \in S\right\}, \tag{1}$$

where $k, \ell \in \mathbb{Z}_+, n \in \mathbb{N}$, $R \in \mathbb{R}^{n\times k}$ and $P \in \mathbb{R}^{n\times \ell}$ are matrices. Denote the columns of matrices R and P by $r_1, \ldots, r_k$ and $p_1, \ldots, p_\ell$, respectively. We allow the possibility that $k = 0$ or $\ell = 0$ (but not both). This general model contains as special cases classical optimization models such as *mixed-integer linear optimization* and *mixed-integer convex optimization*. Given $n \in \mathbb{N}$ and a closed subset $S \subseteq \mathbb{R}^n$ such that $0 \notin S$, a *cut-generating pair* (ψ, π) for S is a pair of functions $\psi, \pi : \mathbb{R}^n \to \mathbb{R}$ such that

$$\sum_{i=1}^{k} \psi(r_i)s_i + \sum_{j=1}^{\ell} \pi(p_j)y_j \geq 1 \tag{2}$$

is a valid inequality (also called a *cut*) for the set $X_S(R, P)$ for every choice of $k, \ell \in \mathbb{Z}_+$ and for all matrices $R \in \mathbb{R}^{n\times k}$ and $P \in \mathbb{R}^{n\times \ell}$. Cut-generating pairs thus provide cuts that separate 0 from the set $X_S(R, P)$. We emphasize that cut-generating pairs depend on n and S and do *not* depend on k, ℓ, R and P. There is a natural partial order on the set of cut generating pairs; namely, $(\psi', \pi') \leq (\psi, \pi)$ if and only if $\psi' \leq \psi$ and $\pi' \leq \pi$. The minimal elements under this partial ordering are called *minimal cut-generating pairs.*

Several deep structural results were obtained by Johnson about minimal cut-generating functions for S when S is a translated lattice, i.e., $S = b + \mathbb{Z}^n$ for some $b \in \mathbb{R}^n \setminus \mathbb{Z}^n$. However, a major drawback is that the theory developed is abstract and difficult to use from a computational perspective. A recent approach has been to restrict attention to a specific class of minimal cut-generating pairs for which we can give computational procedures to compute the values $\psi(r_i)$ and $\pi(p_j)$. We show how this is done when S is a translated lattice intersected with a polyhedron, i.e., $S = (b + \mathbb{Z}^n) \cap Q$ for some vector $b \in \mathbb{R}^n \setminus \mathbb{Z}^n$ and some rational polyhedron Q.

Given such a set $S \subseteq \mathbb{R}^n$, define $W_S := \mathbb{Z}^n \cap \operatorname{lin}(\operatorname{conv}(S))$. A convex set B is called S-free if $\operatorname{int}(B) \cap S = \emptyset$. A *maximal* S-free convex set is an S-free convex set that is inclusion wise maximal. A well-known theorem others). It is known a maximal S-free convex set B containing the origin in its interior is a polyhedron given by

$$B = \{r \in \mathbb{R}^n : a_i \cdot r \leq 1 \ \ i \in I\}. \tag{3}$$

There is an important theorem that says the following: define the following pair of functions associated with B:

$$\psi_B(r) = \max_{i \in I} a_i \cdot r, \quad \pi_B(r) = \inf_{w \in W_S} \psi(r + w) \tag{4}$$

(ψ_B, π_B) is a valid cut-generating pair. Moreover, the pair is "partially" minimal: for every cut-generating pair $(\psi, \pi) \leq (\psi_B, \pi_B)$, we must have $\psi = \psi_B$.

Thus, for every maximal S-free convex set B, (**??**) gives formulas to compute with the corresponding cut-generating pair (ψ_B, π_B). However, because of the partial minimality of (ψ_B, pi_B), it may be the case that there exists a pair (ψ, π) with $\pi \leq \pi_B$ and $\pi(r) < \pi_B(r)$ for some $r \in \mathbb{R}^n$. There are several interesting questions of research. For example, several participants of the workshop were interested on the classification of maximal S-free convex sets. Another question is: Let $S = (b + \mathbb{Z}^n) \cap Q$ with $b \in \mathbb{R}^n \setminus \mathbb{Z}^n$ and a rational polyhedron Q. Given a maximal S-free convex set B (**??**), how can we decide if (ψ_B, π_B) is minimal? We had several discussions about this and during the problem session, the topic was deeply analyzed.

Linear and semidefinite programming are two core optimization technologies with many important applications in mathematics, engineering, and business. An extended formulation is a higher dimensional description of the problem utilizing additional auxiliary variables. Sebastian Pokutta (Georgia Tech, USA) presented a rather interesting survey about extended formulations. The main goal was understand how to reduce the number of required inequalities in a linear programming formulation by representing a given optimization problem in slightly higher dimensional space or ruling out the existence of such formulations. Recently, extended formulations gained significant interest due to fundamental questions in optimization and complexity theory that are closely related to the notion of extended formulations. In fact, extended formulations provide an alternative measure of 'complexity', which is independent of P vs. NP: we count the *number* of required inequalities and the encoding of the coefficients is disregarded. This distinctive criterion makes extended formulations very attractive as the obtained statements are not subject to any complexity theoretic assumptions and it has been argued that the resulting notion of complexity is more in line with how we solve linear programs. Moreover, this notion of complexity might also provide supporting evidence for several conjectures in complexity theory.

More formally, our setup will be the following. Let $P = \{x \mid Ax \leq b\} \subseteq \mathbb{R}^n$ be a polytope representing a combinatorial optimization problem of interest. A polytope $Q = \{x \mid Ex \leq d\} \subseteq \mathbb{R}^m$ with $m \geq n$ is called an *extension of* Q if there exists a linear map π with $P = \pi(Q)$. The smallest number of inequalities required in any extension of P is called the *extension complexity* $\operatorname{xc}(P)$ *of* P. Any extension Q can be used as a surrogate to optimize over P and thus we are interested in finding the smallest possible extension. We therefore ask: *What is the smallest number of inequalities required*
in any extension Q *of* P*?*

Put differently, we aim for determining the extension complexity of P. In many cases using an extended formulation can lead to an exponential saving in terms of the number of inequalities, i.e., a polytope P with an *exponential number* of inequalities in the description $Ax \leq b$ can be expressed in slightly higher dimensional space with a *polynomial* number of inequalities, allowing for efficient optimization over P via linear programming (provided the coefficients are small). Examples include the Spanning Tree Polytope as well as the extended formulations for the regular polygon, which can be used to approximate the second-order cone efficiently. In several other important (and surprising cases), such as e.g., the Traveling Salesman Polytope and Matching Polytope it can be shown that such compact formulations cannot exist.

The theory also naturally extends to approximate formulations and many surprising examples have been recently obtained. For example, it was shown that the MaxCut Problem cannot be approximated better than $1/2$ with a polynomial size linear program. Also, the VertexCover Problem cannot be approximated better than a factor of 2 using a polynomial size linear program.

Pokutta provided an introduction to extended formulations and survey many of the aforementioned results in extended formulations, both in the linear and the semidefinite setting. He explained a reduction framework for establishing upper and lower bounds for the size of exact and approximate LP and SDP formulations. This framework allows for surprisingly simple and convenient analysis without relying on any heavy machinery, making extended formulations very accessible without requiring any in-depth prior knowledge of those results establishing the base hardness. I will conclude with various open problems both in the exact and approximate as well as linear and semidefinite case.

Several other speakers touched on closely related themes for Speaker Levent Tunçel (University of Waterloo, Canada) considered also extended formulations some and discussed the performance of some of the strongest lift-and-project operators in computing the convex hull of integral points inside the given elementary polytope. The discussion included an analysis of the number of major iterations required, as well as an analysis of the changes in the integrality gaps throughout these major iterations. Juan Pablo Vielma (MIT, USA) consider another exciting topic in the challeng

We consider strong Mixed Integer Programming (MIP) formulations for a disjunctive constraint of the form

$$x \in \bigcup_{i=1}^{n} C_i \tag{5}$$

where $\{C_i\}_{i=1}^n \subseteq \mathbb{R}^d$ is a finite family of compact convex sets. MIP formulations for (**??**) can can be divided into two classes depending on their strength and types of auxiliary variables. The first class corresponds to extended formulations that use both 0-1 and continuous auxiliary variables. Standard versions of such extended formulations have sizes that are linear on appropriate size descriptions of the convex sets (e.g. number of linear, quadratic or conic constraints) and have continuous relaxations with extreme points that naturally satisfy the integrality constraints on the 0-1 variables (such formulations are usually denoted *ideal* and are as strong as possible). Extended formulations for polyhedral sets have been introduced by Balas, Jeroslow and Lowe, for conic representable sets by Ben-Tal, Helton, Nemirovski and Nie and for sets described through non-linear inequalities by Ceria, Merhotra, Soares and Stubs. The second class corresponds to non-extended formulations that only use the 0-1 variables that are strictly necessary for a valid formulation. Standard versions of such non-extended formulations are also linear sized, but are often significantly weaker than their extended counterparts. Non-extended formulations include big-M type constraints and ad-hoc formulations for specially structured polyhedral sets

A common feature of both classes is the use of n 0-1 variables that are constrained to add up to one. However, in the polyhedral setting different uses of 0-1 variables can lead to non-extended formulations that are ideal and smaller than the smallest extended counterpart. This allows such formulations to provide a significant computational advantage for disjunctive constraints related to the modeling of piecewise-linear functions. In this talk we describe a systematic geometric procedure to construct such non-extended formulations with a flexible use of 0-1 variables in an attempt to explain and expand on the success of the formulations from prior work with S. Ahmed and G. L. Nemhauser. This procedure is based on an embedding of the disjunctive constraint into a higher dimensional space and leads to several theoretical questions concerning the complexity of unions of polyhedra and the mixed basic semi-algebraic representability of unions of convex basic semi-algebraic sets.

Oktay Gunluk (IBM Research) discussed how to derive cutting planes from extended LP formulations Given a mixed-integer set defined by linear inequalities and integrality requirements on some of the variables, Gunluk considered extended formulations of its continuous (LP) relaxation and studied the effect of adding cutting planes in the extended space. In terms of optimization, extended LP formulations do not lead to better bounds as their projection onto the original space is precisely the original LP relaxation. However, adding cutting planes in the extended space can lead to stronger bounds. He showed that for every 0-1 mixed-integer set with n integer and k continuous variables, there is an extended LP formulation with $2n + k-$ variables whose elementary split closure is integral. The proof is constructive but it requires an inner description of the LP relaxation. He discussed extending this idea to general mixed-integer sets and construct the best extended LP formulation for such sets with respect to lattice-free cuts. We also present computational results on the two-row continuous group relaxation showing the strength of cutting planes derived from extended LP formulations.

Another algorithmic structure that was discussed is the use of group theory. The infinite group problem was introduced 42 years ago by Ralph Gomory and Ellis Johnson in their groundbreaking papers titled "Some continuous functions related to corner polyhedra I, II". The technique, investigating strong relaxations of integer linear programs by convexity in a function space, has at times been dismissed as "esoteric". Now we recognize the infinite group problem as a technique which was decades ahead of its time, providing the first "cut-generating function" approach to integer programming. It may be the key to today's pressing need for stronger, "multi-row" cutting plane approaches. Matthias Köppe (Univ. of California, Davis, USA) surveyed the recent progress on the problem, focusing on algorithmic aspects, such as the automatic extremality test for cut generating functions in the Gomory-Johnson model, its implementation in software, and ongoing work on automatic discovery and proof of cutting plane theorems in the Gomory-Johnson model.

It is well-known that in many problems the constraints are sparse (few non-zero entries). Therefore it makes sense to exploit sparsity in computation. Daniel Bienstock , (Columbia University) discussed how to exploit structured sparsity in mixed-integer polynomial optimization Many ideas in (continuous) polynomial optimization algorithms make use of the structural sparsity of the intersection graph of the constraints (e.g., Waki et al, Lasserre et al). Often this leads to, e.g., sum-of-squares or semidefinite relaxations of the original problem, whose solution is made more efficient by leveraging the sparsity; however concrete convergence results are scarce. Bienstock described linear programming approximations to mixed-integer polynomial optimization problems where the intersection graph of the constraints has fixed tree-width. The LP formulations, given epsilon ¿ 0, are polynomially large in the problem data and in ϵ^{-1}, and provably attain epsilon-optimality and feasibility guarantee. As a consequence he showed how to obtain an LP-based polynomial-time approximation algorithm for several problems arising in energy markets.

Also on the theme of using sparsity, Santanu Dey (Georgia Tech, USA) presented an analysis of sparse cutting-planes for sparse MILPs with applications to stochastic MILPs. Numerous families of cutting-planes have been studied for mixed integer linear programs (MILPs), significantly lesser understanding has been obtained on the very important question of cutting-plane selection from a theoretical perspective. State-of-the-art MILP solvers bias the selection of cutting-planes towards sparse cuts: This is a natural choice since solving a MILP involves solving many linear programs (LP) and LP solvers can take advantage of sparse constraint matrices. In a recent work (Dey, Molinaro and Wang) presented a geometric analysis of the quality of sparse cutting-planes as a function of the number of vertices of the integer hull, the dimension of the polytope and the level of sparsity. Dey discussed the question of understanding the strength of sparse cutting-planes using completely different techniques, so that we are also able to incorporate the information that most real-life MILP formulations have sparse constraint matrices.

Several speakers and discussions concentrated on trying to use very concrete structure in the format of problems. E.g, Kurt Anstreicher (University of Iowa) spoke about the *trust-region subproblem* which is the problem of minimizing a (possibly nonconvex) quadratic objective over an n-dimensional sphere. He discussed considered a generalization of the trust-region subproblem that adds linear inequality constraints as well as hollows, or excluded regions, corresponding to the complements of convex ellipsoids that are contained in the sphere. He proved that this highly nonconvex problem can be solved exactly via semidefinite programming with added RLT and SOC-RLT constraints so long as none of the linear constraints or hollows intersect one another within the sphere.

Touching on another rather famous structure Alejandro Torrielo (Georgia Tech, USA) discussed relaxations for a Dynamic Knapsack Problem We consider a dynamic version of the classical knapsack problem with the following formulation. Let $N := \{1, \ldots, n\}$ be a set of items. For each item $i \in N$ we have a non-negative, independent random variable A_i with known distribution representing its size, and a deterministic value $c_i > 0$. We have a knapsack of deterministic capacity $b > 0$, and we would like to maximize the expected total value of inserted items. An item's size is realized when we choose to insert it, and we receive its value only if the knapsack's remaining capacity is greater than or equal to the realized size. Given any remaining capacity $s \in [0, b]$, we may choose to insert any available item, and the decision is irrevocable. If the insertion is unsuccessful, i.e. the realized size is greater than the remaining capacity, the process terminates.

This model and others like it have applications in scheduling, equipment replacement and machine learning, to name a few examples. They also generally reflect some trends in optimization research and its various application, which have focused attention on models in which uncertain data is not revealed at once after an initial decision stage, but rather is dynamically revealed over time based on the decision maker's choices.

The deterministic knapsack is a special case, so this problem is NP-hard, and some variants are known to be PSPACE-hard. Because the decision maker can choose any item to insert based on remaining capacity, a solution is not simply a subset of items, but rather a policy that prescribes what item to insert under all possible circumstances. Research on the model has therefore studied heuristic policies and tractable relaxations. Our focus is mostly on the latter, deriving mathematical programming relaxations that can be solved efficiently, and which can be used to design high-quality heuristics. Specifically:

1. We introduce a semi-infinite relaxation for the problem under arbitrary item size distributions, based on an affine value function approximation of the linear programming encoding of the problem's dynamic program. We show that the number of constraints in this relaxation is at worst countably infinite, and is polynomial in the input for distributions with finite support.

2. When item sizes have integer support, we show that non-parametric value function approximation gives the strongest known relaxation from the literature, which has pseudo-polynomially many variables and constraints.

3. We theoretically and empirically compare these relaxations to others from the literature and show that both are quite tight. In particular, our new relaxation is notably tighter than a variety of benchmarks and compares favorably to the theoretically stronger pseudo-polynomial relaxation when this latter bound can be computed.

Time permitting, we also discuss future work and open questions motivated by our results, including the theoretical worst-case gap of our new relaxation, the possible strengthening of our relaxations, their asymptotic behavior as the number of items grows, and others.

2.3 Applications

Fascinating applications motivate the mathematical methods of discrete optimization methods.

Speaker Dorit Hochbaum (Univ. of California, Berkeley, USA) presented a discrete optimization model for clustering data which combines two criteria: Given a collection of objects with pairwise similarity measure, the problem is to find a cluster that is as dissimilar as possible from the complement, while having as much similarity as possible within the cluster. The two objectives are combined either as a ratio or with linear weights. The ratio problem, and its linear weighted version, are solved by a combinatorial algorithm within the complexity of a single network minimum s,t-cut algorithm. We call this problem "the normalized cut prime" (NC') as it is closed related to the NP-hard problem of normalized cut. The relationship of NC' to normalized cut is generalized to a problem we call "q-normalized cut". It is shown that the spectral method that solves for the Fielder eigenvector of a related matrix is a continuous relaxation of the problem. In contrast, the generalization of the combinatorial algorithm solves a discrete problem resulting from a relaxation of a single sum constraint. Hochbaum discussed the relationship between these two relaxations and explained a number of advantages for the combinatorial algorithm. These advantages include a better approximation, in practice, of the normalized cut objective for image segmentation benchmark problems. There are applications of NC', as a supervised machine learning technique, to data mining, and it compares favorably to leading machine learning techniques on datasets selected from data mining benchmark.

Justo Puerto from (Univ. de Sevilla, Spain) discussed another family of applications called the k-sum and ordered median combinatorial optimization problems. In his talk, he addressed the continuous, integer and combinatorial k-sum and ordered median optimization problems. He discussed the analysis of different reformulations of these problems that allow to solve them through the minimization of mini-sum optimization problems. This approach provides a general tool for solving ordered median optimization problems and improves the complexity bounds of many ad-hoc algorithms previously developed in the literature for particular versions of these problems.

Roger Z. Rios-Mercado (Univ. Autónoma de Nuevo León) talked about a rather interesting problem related to the creation of districts (a problem that appears in product distribution and in political districting). Districting or territory design involves locational decisions where a given set of basic or geographic units must be partitioned so as to optimize some performance measure subject to pre-specified planning requirements. Typical criteria usually sought are territory compactness, connectivity, balancing, similarity with existing plan, etc. Depending on the particular application, different models or dispersion measure can be used. Rios-Mercado gave an overview of the main elements involving districting decisions, and present some of the models, and solution algorithms (exact and heuristic) that have been developed for particular districting applications.

Francisco Zaragoza ((UAM Azcapotzalco, Mexico) spoke about another application in scheduling problems that can be modeled with the Traveling Repairman Problem on a Line with Unit Time Windows. More precisely, let $G = (V, E)$ be a graph and $r \in V$. For each $e \in E$, let $\ell_e > 0$ be the length of e. For each $v \in V$, a time window $[a_v, b_v]$ is given. A repairman starts in vertex r at time $t = 0$ and moves trough the edges of G at unit speed. The Traveling Repairman Problem consists of finding a route for the repairman that maximizes the number of vertices visited during their time windows. This problem is known to be NP-hard even when G is a tree and each time window has unit length (Frederickson and Wittman, 2012) or when G is a path and time windows are arbitrary (Tsitsiklis, 1992). The complexity of the remaining case, that is when G is a path and all time windows are unitary, is still open. Much work has been done in this case: there are approximation algorithms with guarantees 8 and $4 + \epsilon$ in quadratic time (Bar-Yehuda, Even, Shahar, 2005) and 3 in quartic time (Frederickson and Wittman, 2012). The algorithm with guarantee 8 has been improved, to get a guarantee of 4 in quadratic time (López, Pérez, Urban, and Zaragoza, 2014). All these algorithms use dynamic programming as the main tool. Zaragoza and collaborators have improved the analysis of this algorithm to show that it guarantees less than 3. The main tool was setting up a linear program that describes the possible outcomes of the dynamic program, and solving it to find the worst possible outcome.

3 Outcome of the Meeting

The workshop was quite successful. We had a total of 42 participants from institutions of Canada, Mexico and the USA. The themes discussed at the workshop were an excellent representation of the very latest trends in research; for example, there were several presentations devoted to discrete-optimization problems using non-linear constraints. The presentations each day touched on more than one topic to provide variety (see schedule), It is worth stressing that several joint projects were started as a result of the meeting and the problem sessions were extremely active. Several collaborations were started during the meeting and we know they will lead to interesting outcomes. Many of the participants expressed their thanks to the wonderful staff support we received, from the early organization (e.g., sending invitations) to the moment of the conferences where CMO staff took extremely good care of us!! Thanks!!