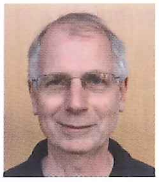  

William Cook *University of Waterloo, Canada*
Martin Grötschel *Zuse Institute and TU Berlin, Germany*
Alexander Schrijver *CWI and University of Amsterdam, Netherlands*

The Year Combinatorics Blossomed

One summer in the mid 1980s, Jack Edmonds stopped by the Research Institute for Discrete Mathematics in Bonn for an extended visit. As usual, the institute administrator asked Professor Edmonds for a curriculum vitae to complete the university paperwork. The conversation took place in the library, so Edmonds pulled from a nearby shelf a text in combinatorial optimization: *"Here is my CV."*

And he was right! This year marks the 50th anniversary of the publication of two papers by Edmonds that, together with his own follow-up work, have come to define much of the field, including theory, complexity, and application. We thought it fitting to write a few modest words on the profound impact of these papers. This short article will not go into the subject in any detail, but, for that, you can check out any book with "combinatorial optimization" in the title — Edmonds's work will fill the pages from cover to cover.

Linear Programming (LP)

Although our focus is Edmonds, to put his contributions into context we have to first go back to the 1940s and the introduction of the linear-programming model by George Dantzig. Indeed, in the first of the two big manuscripts in 1965, Edmonds [5] writes the following.

This paper is based on investigations begun with G. B. Dantzig while at the RAND Combinatorial Symposium during the summer of 1961.

For a brief time, the two great mathematicians worked side by side, and their discussions set Edmonds on the course towards developing one of the most important settings for Dantzig's optimization theory.

Much has been written about linear programming, including several hundred texts bearing the title. Dantzig's creation of the model and the simplex algorithm for its solution is rightly viewed as one of the greatest contributions of applied mathematics in the past century. For our purposes, it will suffice to give the briefest of descriptions.

Every LP problem can be formulated as the task to minimize (or maximize) a linear function subject to linear equality or inequality constraints and non-negative values for the variables. That is, a model with n variables and m constraints can have the form

$$\begin{aligned}
&\text{minimize } c_1x_1 + c_2x_2 + \ldots + c_nx_n\\
&\text{subject to}\\
&a_{11}x_1 + a_{12}x_2 + \ldots + a_{1n}x_n \geq b_1\\
&a_{21}x_1 + a_{22}x_2 + \ldots + a_{2n}x_n \geq b_2\\
&\ldots\\
&a_{m1}x_1 + a_{m2}x_2 + \ldots + a_{mn}x_n \geq b_m\\
&x_1 \geq 0, x_2 \geq 0, \ldots, x_n \geq 0.
\end{aligned}$$

Here, the x_j values are the unknowns, while the c_j, a_{ij}, and b_i values are given as part of the problem. In brief matrix notation, the LP model stated above reads $\min c^T x \text{ s.t. } Ax \geq b, x \geq 0$.

The economic interpretation of the general model is that the x_j variables represent decisions, such as the quantity of certain items to purchase; the c_j values are the costs of one unit of each item; and the constraints capture requirements on the portfolio of items that are purchased. The canonical textbook example, and one that Dantzig himself considered in early tests of the simplex algorithm, is the *diet problem*, where there are n food items that can be purchased and m nutritional requirements, such as the minimum number of calories, grams of protein, etc., that must be included in the daily selection of food. The LP solution provides the cheapest way to keep a person on his or her feet for a day.

The general model is a simple one, and that is one of the reasons for its success: pretty much any industry you can name makes use of linear-programming software to guide their decision making.

The mathematical elegance of linear programming is tied to the fact that to each LP problem we can associate another problem called its *dual*. The dual LP problem is obtained by turning the model on its side, having a dual variable y_i for each original constraint and a dual constraint for each of the original variables:

maximize $b_1y_1 + b_2y_2 + \ldots + b_my_m$
subject to
$a_{11}y_1 + a_{21}y_2 + \ldots + a_{m1}y_m \le c_1$
$a_{12}y_1 + a_{22}y_2 + \ldots + a_{m2}y_m \le c_2$
...
$a_{1n}y_1 + a_{2n}y_2 + \ldots + a_{mn}y_n \le c_n$
$y_1 \ge 0, y_2 \ge 0, \ldots, y_m \ge 0.$

In matrix notation, the meaning of "turning the model on its side" becomes even more visible: max b^Ty s.t. $A^Ty \le c, y \ge 0$.

To keep the names straight, Dantzig's father, Tobias, proposed the original LP problem be called the *primal* problem. So primal and dual.

A simple result is that for any $x_1, \ldots, x_n$ values satisfying the primal LP constraints and any $y_1, \ldots, y_m$ values satisfying the dual LP constraints, we have $c_1x_1 + \ldots + c_nx_n$ cannot be less than $b_1y_1 + \ldots + b_my_m$. Indeed, $c^Tx = x^Tc \ge x^TA^Ty = (Ax)^Ty \ge b^Ty$. So any candidate solution to the dual gives a bound on how small we can make the primal objective, and, vice versa, any candidate solution to the primal gives a bound on how large we can make the dual objective. A deeper result, called the *LP Duality Theorem*, is that an optimal solution $x_1^*, \ldots, x_n^*$ to the primal problem and an optimal solution $y_1^*, \ldots, y_m^*$ to the dual problem will have equal objective values, that is,

$$c_1x_1^* + \ldots + c_nx_1^* = b_1y_1^* + \ldots + b_my_m^*.$$

If you like mathematics, then you have to love the Duality Theorem. The equation gives a concise way to prove to any sceptic that you have in hand an optimal solution to a given LP model: you simply display a dual solution that gives the same objective value. Dantzig's simplex method proceeds by simultaneously solving the primal and dual LP problems, each solution providing an optimality certificate for the other.

Figure 1 Jack Edmonds, September 2014

Combinatorial Min-Max Theorems

It didn't take long for mathematicians to realize that LP duality was too pretty to just sit on a shelf. In an incredibly active span of years in the mid-1950s, researchers rapidly expanded the reach of LP theory and Dantzig's algorithm.

Alan Hoffman, the long-time master of combinatorics and linear algebra wrote the following in a memoir [10].

It dawned on me (and on Gale, Kuhn, Heller, Tucker, Dantzig, Ford, Fulkerson, Kruskal, Tompkins and others) that you could prove combinatorial theorems by using the theory of linear inequalities! And you could try to discover and understand the machinery that ensured integrality of optimal solutions, because integrality was needed to make the theorems combinatorial. This was incredible. I cannot overstate how this discovery boosted my morale.

Typical of the morale boosters was the LP-based proof of König's Theorem in graph theory, stating that in a bipartite graph, the minimum number of nodes that together meet all edges is equal to the maximum number edges, no two of which meet at a node.[1] The theorem is illustrated in **Figure 2**, where there are three highlighted nodes and three highlighted edges.

To put König's Theorem into the LP setting, we have primal variables $x_1, \ldots, x_n$, one for each of the n nodes in the graph, and a constraint for each of the m edges, stating that the sum of the variables for the two ends of the edge must be at least one.

—
1 A graph is called bipartite if its nodes can be colored red and blue so that every edge has one red end and one blue end.

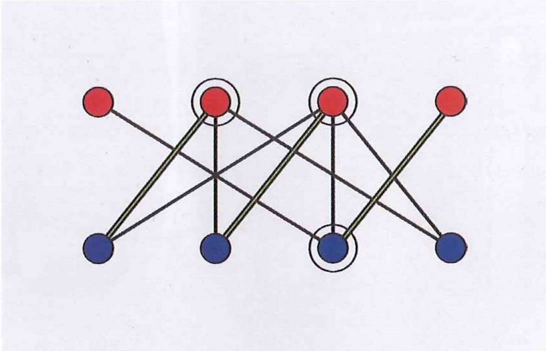

Figure 2 König's Theorem example

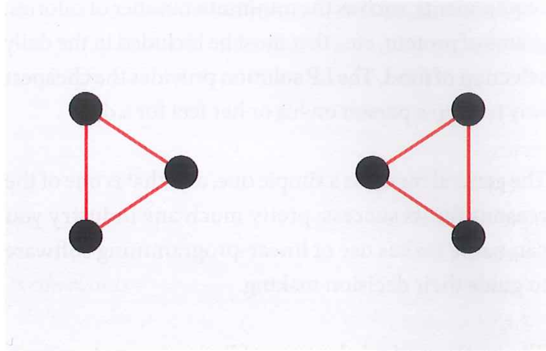

Figure 3 Red edges carry the value 1/2 in the LP solution.

minimize $x_1 + x_2 + \ldots + x_n$
subject to
$x_1 + x_j \geq 1$, for each edge (i, j)
$x_1 \geq 0, x_2 \geq 0, \ldots, x_n \geq 0.$

The dual LP model flips this around; we have variables $y_1, \ldots, y_m$, one for each edge, and a constraint for each node, stating that the sum of the variables for the edges meeting that node can be at most one.

maximize $y_1 + y_2 + \ldots + y_m$
subject to
$\sum(y_j : \text{edge } j \text{ meets node } i) \leq 1,$ for each node i
$y_1 \geq 0, y_2 \geq 0, \ldots, y_m \geq 0.$

Although LP models will in general have fractional-valued solutions, Hoffman and company found direct arguments showing that, for any bipartite graph, this primal and dual pair have optimal solutions where all variables have value either zero or one. These zero-one solutions pick out sets of nodes and edges, corresponding to the variables that carry the value one. Now, for these sets, the LP constraints are precisely the conditions we imposed on the selection of nodes and edges in König's Theorem. Thus, the result of König follows from the Duality Theorem.

The "1950s gang" set out to apply this LP-duality scheme to every combinatorial problem they could find. And they had success upon success. But some models resisted, such as the traveling salesman problem (TSP) and the matching problem in general graphs. The difficulty was that there are instances of these problems where the LP models have optimal solutions only with some or all variables taking on fractional values.

Paths, Trees, and Flowers

The LP scheme was powerful, but it could not be coaxed into a general theory for combinatorial problems. Not, that is, until Edmonds's big year.

The model Edmonds attacked in 1965 was the perfect matching problem. A *perfect matching* in a graph is a set of edges that meet every node exactly once. Given a cost c_i associated with each edge i, the problem is to find a perfect matching of minimum total cost. Thus, we need a zero-one valued solution to the model

minimize $c_1x_1 + \ldots + c_mx_m$
subject to
$\sum(x_j : \text{edge } j \text{ meets node } i) = 1,$ for each node i
$x_1 \geq 0, x_2 \geq 0, \ldots, x_m \geq 0$

where the variables $x_1, \ldots, x_m$ correspond to the m edges in the graph.

The perfect matching problem includes, as a special case, the geometric challenge of pairing up points in the plane, so that the sum of the lengths of the lines joining the pairs is as small as possible. In this geometric setting, we can see easily what goes wrong with the LP approach. Consider an example consisting of two clusters, each with three points. Any perfect matching must include an edge joining a point in one cluster to a point in the other cluster, but the LP solution to the above model will instead create two triangles of edges, each carrying the value of 1/2, as we illustrate in **Figure 3**. And there is no way to avoid such a non-matching solution: for that particular set of points, the 1/2-values form the unique optimal solution to the LP model.

These bad LP solutions can be described in geometric terms; not in the 2-dimensional space of the points we want to match, but rather in the space where we have a dimension for every edge of the graph. Indeed, the set of candidate solutions to a LP model together form a geometric object called a *polyhedron*. Think of a Platonic solid, like a dodecahedron, but in high-dimensional space. The linear inequality constraints in the LP model form the sides, or *faces*, of the polyhedron. A polyhedron is a *convex* set, that is, if you take any two points u and v in a polyhedron, then the entire line segment joining u and v is also in the polyhedron. The *vertices* of a polyhedron are the corner points, that is, those points p in the polyhedron for which there do not exist distinct points u and v in the polyhedron such that p is on the (u, v) line segment. The vertices are special: an optimal solution to an LP model can always be found among its vertices, and, for any vertex, there is a way to set the costs c_i of the variables so that the vertex is the unique optimal solution.

In our LP model, every perfect matching determines a vertex, by setting $x_i = 1$ if edge i is in the matching and otherwise setting $x_i = 0$. That is good. But 1/2-valued solutions, like in our 6-node example, are also vertices. That is bad.

What we want is a polyhedron where every perfect matching is a vertex, and these are the only vertices. Such a polyhedron always exists. Indeed, at the turn of the 20th century, Hermann Minkowski showed that for any finite set of points X, if we let P be the smallest convex set containing X, then P is a polyhedron. Minkowski's set P is called the *convex hull* of X.

Minkowski's theorem is well known to students of linear programming, but it says only that an LP model for perfect matchings exists. It does not say how we should find the inequalities to use as LP constraints. And if even if we can find them, the form of the inequalities might make them too complex or too numerous to use in any nice theory for matchings. These are the difficulties Edmonds handled. In so doing, he created a road map for the potential solution to any problem in combinatorial optimization.

The matching problem is one of the oldest in graph theory, and Edmonds had at his disposal results dating back to the late 1800s. Much of the theory points towards a central role for odd-cardinality subsets of nodes. Indeed, if a subset S contains an odd number of nodes, then a perfect matching of the graph must include at least one edge joining a node in S to a node not in S. In other words, every perfect matching satisfies the linear inequality

$$\sum(x_e : e \text{ has one end in } S) \geq 1$$

Edmonds calls these constraints blossom inequalities. His theorem is that adding these inequalities, for every odd set S, gives the convex hull of perfect matchings. That is, not only do the inequalities cut off all half-integer solutions, such as our two red triangles, they also do not introduce any new vertices. Remarkable!

Edmonds's proof is via an efficient algorithm that constructs a perfect matching and a corresponding dual solution that together satisfy the LP-duality equation. For geometric instances, the dual solution can be viewed as a set of nested regions trapping in odd sets of points, as we illustrate in **Figure 4** with an optimal matching of 50 points.[2] Note that although there are an impossibly large number of blossom inequalities on 50 points, the dual solution has only a modest number of variables that take on positive values. This is a direct consequence of Edmonds's algorithm, showing that we need not fear convex hulls having many faces, as long as we understand well the inequalities that produce those faces.

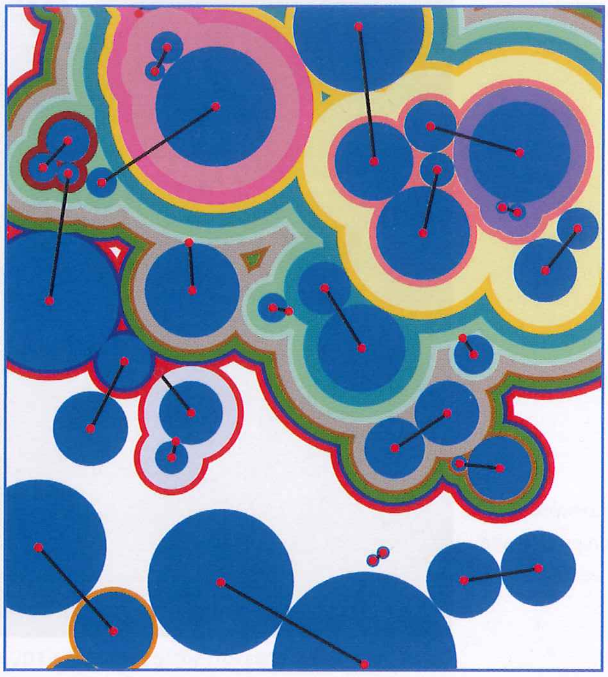

Figure 4 Optimal matching and Edmonds dual solution.

"It was my first glimpse of heaven" as Edmonds would later state [7].

The field that has grown up around Edmonds's approach is called *polyhedral combinatorics*, where one takes a combinatorial problem and aims to create both efficient solution algorithms and pretty theorems, such as König's min-max equation. Alan Hoffman [10] writes the following.

A classic mathematician's joke is that the first time you use a new technique it is a trick, the second time it's a method, the third time a topic. Clearly, polyhedral combinatorics has become a subject, which engages some of the world's outstanding combinatorial mathematicians in a big fraction of their research.

Edmonds's work has certainly had a unifying effect on the optimization side of combinatorics. Vašek Chvátal's [1] slogan

combinatorics = number theory + linear programming

sums up nicely the overall approach.

—

2 This visualization technique is due to Michael Jünger and William Pulleyblank [11].

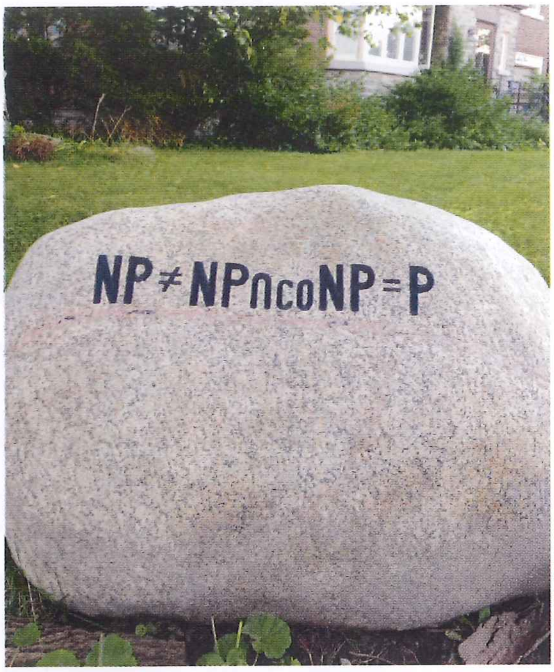

Figure 5 The Complexity Rock at Jack Edmonds's house in Ontario, Canada.

Good Complexity, Bad Complexity

We used above the word "efficient" to distinguish Edmonds's algorithm from the brute-force approach of listing each perfect matching and selecting the cheapest. This is a point Edmonds made clear in his paper [5], when describing another of his matching algorithms.

I am claiming, as a mathematical result, the existence of "a good" algorithm for finding a maximum cardinality matching in a graph. There is an obvious finite algorithm, but that algorithm increases in difficulty exponentially with the size of the graph. It is by no means obvious whether "or not" there exists an algorithm whose difficulty increases only algebraically with the size of the graph.

The algorithms Edmonds calls "good" are now called *polynomial-time algorithms*, or the class $\mathcal{P}$ for short. A lengthy discussion in his paper—often overlooked in computer science articles on the subject— became the basis for the development of much of modern computational complexity theory, including the fundamental $\mathcal{P}$ versus $\mathcal{NP}$ question that is one of the million-dollar Clay Prize problems.

Applications Want Solutions

Several years after his success with matchings, Edmonds became convinced that the TSP was beyond the reach of a polynomial-time solution [6]: "I conjecture that there is no good algorithm for the traveling salesman problem." With the development of complexity theory by Stephen Cook and Richard Karp, many researchers would today also make this conjecture. It is equivalent to the statement that $\mathcal{P} \neq \mathcal{NP}$, since the TSP is in the class of $\mathcal{NP}$-hard problems, like many other combinatorial models.

It is important to note, however, that the notion of $\mathcal{NP}$ hardness refers to the possibility of bad asymptotic behavior of a problem class. When an $\mathcal{NP}$-hard problem arises in an application, what we need to solve are specific, finitely-sized examples; complexity theory should not deter us from attacking the problem with the mathematical tools at our disposal. Edmonds himself [5] wrote the following when he introduced the notion of good algorithms.

It would be unfortunate for any rigid criterion to inhibit the practical development of algorithms which are either not known or known not to conform nicely to the criterion.

Indeed, it was Edmonds's matching work that prompted a broad study of practical LP-based methods for $\mathcal{NP}$-hard problems such as the TSP.

In this line of work, partial descriptions of the convex hulls are utilized to obtain, via LP duality, strong bounds on the value of the cost of a best-possible solution. For example, although we do not know all of the inequalities needed to obtain the convex hull of TSP solutions, we know enough of them to be able to produce strong statements of the form: "No tour through these points can have length less than X kilometers." This mechanism can then be utilized in an enumerative process to locate the optimal tour and to prove it is the shortest possible, allowing computer implementations to solve routinely TSP examples with 1,000 or more cities.

The overall procedure, known as the *cutting-plane method or branch-and-cut*, is a powerful tool for the solution of a wide range of models arising in industry and commerce. It has its roots in work by Dantzig and colleagues in the early 1950s [3], but the rapid advancements began only in the years following Edmonds's papers. In fact, today the LP-based branch-and-cut procedure is the corner stone of almost all commercial optimization software packages, and there is almost no product or service in the world where this methodology has not contributed to its design, manufacturing or delivery.

Optimization ≡ Separation

After his work on matchings, Edmonds knew that he had a powerful general framework on his hands. Here is a remark he made in 1964 [8].

For the traveling salesman problem, the vertices of the associated polyhedron have a simple characterization despite their number—so might the bounding inequalities have a simple characterization despite their number. At *least we should hope they have, because finding a really good traveling salesman algorithm is undoubtedly equivalent to finding such a characterization.*

The thesis of Edmonds was clear: the existence of polynomial-time algorithms goes hand-in-hand with polyhedral characterizations.

An awkward point in the study of the complexity of algorithms, however, was that the simplex method itself, that stalwart of efficiency and practicality, was not known to be a good algorithm in the sense of Edmonds.

It remains an open problem to find a *good* simplex algorithm, but linear programming itself did eventually fall under Edmonds's umbrella. Indeed, the most widely circulated news event in the history of mathematical optimization occurred in the summer of 1979, when Leonid Khachiyan published a polynomial-time algorithm for solving LP problems. The story was covered on the front page of the *New York Times* and in other newspapers around the world. Part of the excitement, in that Cold War era, was that Khachiyan's work did not make use of the simplex algorithm, adopting instead the ellipsoid method for convex programming developed by Naum Shor, David Yudin and Arkadi Nemirovski in the Soviet Union.

Claims in the media that Khachiyan had laid to rest the venerable algorithm of Dantzig, as well as solving the TSP along the way, were wildly off base. The ellipsoid LP method did not prove to be viable in practice for the solution of large-scale models.

Ellipsoids did, however, have a great impact on the theory of algorithms. The precise result, known as *optimization ≡ separation*, is technical, but it says, roughly, that Edmonds was right again. If we can solve a combinatorial problem in polynomial time, then we have an implicit description of the corresponding convex hull, and, vice versa, if we understand the convex hull then we have a polynomial-time algorithm for the combinatorial problem. This nicely ties together the polyhedral, algorithmic, and complexity components of Edmonds's work.

Suggested Reading

The book [12] is a comprehensive survey of the theory and methods of polyhedral combinatorics and covers, in particular, the contributions of Jack Edmonds to matching theory and beyond in great detail. Several chapters of the part "Discrete Optimization Stories" of the book [9] give accounts of some of the historical developments in combinatorial optimization. The article by William R. Pulleyblank in this book discusses Edmonds's work on matching and polyhedral combinatorics. The book [2], readable for the nonspecialist, outlines the modern solution approach to combinatorial optimization problems using the traveling salesman problem as an example.

The Next Fifty Years

Polyhedral combinatorics and complexity theory are thriving fields, with ever more connections being made to classical areas of mathematics. And on the applied side, techniques based on polyhedra and cutting planes continue to expand their reach into new problem domains and increasingly complex models. The next fifty years should be an exciting time, as the mathematics world continues to digest Edmonds's glimpse of heaven.

References

[1] Chvátal, V. 1973. Edmonds polytopes and a hierarchy of combinatorial problems. Discrete Mathematics 4, 305–337.

[2] Cook, W. 2012. *In Pursuit of the Traveling Salesman: Mathematics at the Limits of Computation*. Princeton University Press, Princeton, New Jersey.

[3] Dantzig, G., R. Fulkerson, S. Johnson. 1954. Solution of a large-scale traveling-salesman problem. Operations Research 2, 393–410.

[4] Edmonds, J. 1965. Maximum matching and a polyhedron with 0,1-vertices. Journal of Research of the National Bureau of Standards 69B, 125–130.

[5] Edmonds, J. 1965. Paths, trees, and flowers. Canadian Journal of Mathematics 17, 449–467.

[6] Edmonds, J. 1967. Optimum branchings. J. Res. Nat. Bur. Stand. Sec. B 71, 233–240.

[7] Edmonds, J. 1991. A glimpse of heaven. J.K. Lenstra et al., eds. *History of Mathematical Programming—A Collection of Personal Reminiscences*. North-Holland, Amsterdam. 32–54.

[8] Gomory, R.E. 1966. The traveling salesman problem. *Proceedings of the IBM Scientific Computing Symposium on Combinatorial Problems*. IBM, White Plains, New York, USA. 93–121.

[9] Grötschel, M. 2012. *Optimization Stories*. Documenta Mathematica, Bielefeld, Germany.

[10] Hoffman, A. 2007. What the transportation problem did for me. Annals of Operations Research 149, 117–120.

[11] Jünger, M., W.R. Pulleyblank. 1993. Geometric duality and combinatorial optimization. S.D. Chatterji, B. Fuchssteiner, U. Kulisch, R. Liedl, eds. *Jahrbuch Überblicke Mathematik*. Vieweg, Braunschweig/Wiesbaden, Germany. 1–24.

[12] Schrijver, A. 2003. *Combinatorial Optimization: Polyhedra and Efficiency*. Springer, Berlin, Germany.