Limitations of Linear Programming as a model of approximate computation

Yannis Moysoglou*

Department of Informatics and Telecommunications, National and Kapodistrian University of Athens, Panepistimiopolis Ilissia, Athens
gmoys@di.uoa.gr

Abstract. Linear programming has proved to be one of the most powerful and widely used tools in algorithm design and especially in the design of approximation algorithms. It has proved its expressive power by modeling diverse types of problems in planning, routing, scheduling, assignment, and design. However there are problems that seem to be very hard for linear programming. More specifically, the capacitated facility location problem (CFL) is an example of an important and well-studied problem for which, while it can be approximated within a constant factor using local search, it is not known to admit efficient relaxation based approximations.

In this thesis we take the direction of exploring the limitation of linear programming. Most of the thesis's results are concerned with linear programming approximability of the capacitated versions of the metric facility location problem such as the capacitated facility location (CFL). We give impossibility results in the hierarchy and in the extended formulations models and we also study another, independent family of relaxations which we call proper.

We show that the relaxations obtained from the natural LP at $\Omega(n)$ levels of the semidefinite Lovász-Schrijver hierarchy for mixed programs, and of the Sherali-Adams hierarchy, the integrality gap is $\Omega(n)$, where n is the number of facilities. Our bounds are asymptotically tight. Then we prove that the standard CFL relaxation enriched with the submodular inequalities of [1] has also an $\Omega(n)$ gap and thus not bounded by any constant. This disproves a long-standing conjecture of [24].

We propose a framework for proving lower bounds on the size of extended formulations. We do so by introducing specific types of extended relaxations that we call *product and distributional relaxations*. Then we show that for every approximate extended formulation of a polytope P, there is a product or distributional relaxation that has the same size and is at least as strong. We provide a methodology for proving lower bounds on the size of approximate product and distributional relaxations and, as an application of our method, we show for CFL an exponential lower bound on the size of a restricted type of extended formulations.

* Dissertation Advisor: Stavros Kolliopoulos, Associate Professor

1 Dissertation Summary

This thesis is naturally divided in three parts regarding the subfield in which the corresponding results reside: One part regarding the quality of solution for the CFL problem obtained by capitalizing on LP hierarchies, a second part, and perhaps the most important contribution of the thesis, regarding the development of a methodology for lower bounding the size of Extended Formulations and applications to the CFL problem, and one smaller part regarding the characterization of the strength of generalized configuration linear programs for CFL. In this overview we first define CFL and give a detailed backround, and then we give an introduction to each topic corresponding to each part: we mention relevant work and we briefly present the exact contribution and results obtained by the research conducted in the context of this thesis.

1.1 Approximating Facility Location

Facility location is one of the most well-studied families of models in combinatorial optimization. In the *uncapacitated facility location* problem *(*UFL*)* we are given a set F of facilities and a set C of clients. We may open facility i by paying its opening cost f_i and we may assign client j to facility i by paying the connection cost c_{ij}. We are asked to open a subset $F' \subseteq F$ of the facilities and assign each client to an open facility. The goal is to minimize the total opening and connection cost. The approximability of UFL was settled by an $O(\log |C|)$-approximation [15], which via a reduction from Set Cover is asymptotically best possible, unless $\mathsf{P} = \mathsf{NP}$ [28]. In *metric* UFL the service costs satisfy the following variant of the triangle inequality: $c_{ij} \leq c_{ij'} + c_{i'j'} + c_{i'j}$ for any $i, i' \in F$ and $j, j' \in C$. This natural special case of UFL is approximable within a constant-factor, and many improved results have been published over the years. In those, LP-based methods, such as filtering, randomized rounding and the primal-dual method have been particularly prominent (see, e.g., [34]). After a long series of papers the currently best approximation ratio for metric UFL is 1.488 [25], while the best known lower bound is 1.463, unless $\mathsf{P} = \mathsf{NP}$ ([33]).

CFL is the generalization of metric UFL where every facility i has a capacity u_i that specifies the maximum number of clients that may be assigned to i. In *uniform* CFL all facilities have the same capacity U. Finding an approximation algorithm for CFL that uses a linear programming lower bound was until recently a notorious open problem. The natural LP relaxations have an unbounded integrality gap and up to the recent breakthrough of [5], the only known $O(1)$-approximation algorithms were based on local search, with the currently best ratios being 5 [8] for the non-uniform and 3 [2] for the uniform case respectively. In the special case where all facility costs are equal, CFL admits an LP-based 5-approximation [24]. Williamson and Shmoys [34], stated the design of a relaxation-based algorithm for CFL as one of the top 10 open problems in approximation algorithms. Very recently, An et al. [5] gave a polynomial-time LP-based 288-approximation algorithm, thus answering the open question of [34]. The LP in [5] has exponential size and is not known to be separable in polynomial time. Therefore the question on the existence of an efficient, compact, linear relaxation for CFL remains open. The series of our results regarding the LP-(in)approximability of CFL can be taken as very strong evidence that such a relaxation does not exists.

The Lower Bound Facility Location (LBFL) is in a sense the opposite problem to CFL. In an LBFL instance every facility i comes with a lower bound b_i which is the minimum number of clients that must be assigned to i if we open it. In *uniform* LBFL all the lower bounds have the same value B. LBFL is even less well-understood than CFL. The first approximation algorithm for the uniform case had a performance guarantee of 448 [32], which has been improved to 82.6 [3]. Both use local search. Interestingly, the LBFL algorithms from [32, 3] both use a CFL algorithm on a suitable instance as a subroutine.

1.2 Lift-and-Project methods and the resulting Hierarchies as model of LP computation

A lot of effort has been devoted to understanding the quality of relaxations of 0-1 polytopes obtained by an iterative lift-and-project procedure. Such procedures define hierarchies of successively stronger relaxations, where valid inequalities are added at each level. At level at most d, where d is the number of variables, all valid inequalities have been added and thus the integer polytope is expressed. Relevant methods include those developed by Balas et al. [7], Lovász and Schrijver [27] (for linear and semidefinite programs), Sherali and Adams [30], Lasserre [22] (for semidefinite programs). See [23] for a comparative discussion. The seminal work of Arora et al. [6], studied integrality gaps of families of relaxations for Vertex Cover, including relaxations in the Lovász-Schrijver (LS) hierarchy. This paper introduced the use of hierarchies as a restricted model of computation for obtaining LP-based hardness of approximation results.

We give impossibility results on arguably the most promising directions for obtaining efficient linear strengthened relaxations for CFL using hierarchies and, in doing so, we answer open problems from the literature.

Our first result of this part of our contribution is that there is an instance with $\Theta(n)$ facilities and $\Theta(n^4)$ clients on which the relaxations produced at $\Omega(n)$ levels when the LS procedure is applied on the natural CFL LP have an integrality gap of $\Omega(n)$. The natural LP has a facility opening variable y_i, for every $i \in F$, and an assignment variable x_{ij}, for every $i \in F$, and client $j \in C$. We also extend the former result to the mixed LS_+ hierarchy. This procedure is the stronger version of LS where one additionally requires that every protection matrix is positive semidefinite. The *mixed* LS_+ *procedure* for a mixed integer program is the version of LS_+ where one lifts only the 0-1 variables and requires that the resulting protection matrix is positive semidefinite (see, [7], [12]). We show that the $\Omega(n)$ gap applies for $\Omega(n)$ rounds of mixed LS_+ as well.

We then show that the LPs obtained from the natural relaxation for CFL at $\Omega(n)$ levels of the stronger SA hierarchy have a gap of $\Omega(n)$ on the same family of instances used for the LS result, with $|F| = \Theta(n)$ and $|C| = \Theta(n^4)$, giving the second contribution of this part. This result answers the questions of [26] and [4] stated above as far as the natural LP is concerned. Our bound is asymptotically tight since the relaxation obtained at every level of the SA hierarchy is at least as strong as the one obtained at the same level of LS . We use a variation of the *local-to-global* method which was implicit in [6] for local-constraint relaxations and was then extended to the SA hierarchy in [13]. From a qualitative aspect, we give the first, to our knowledge, hierarchy bounds for a relaxation where variables have more than one type of semantics, namely the facility opening and the client assignment type. Compare this, for example, with the Knapsack and Max Cut LPs that contain each one type of variable.

Our third contribution in this part is that the *submodular* inequalities introduced in [1] for CFL fail to reduce the gap of the classic relaxation to constant. These constraints generalize the flow-cover inequalities for CFL. Thus we disprove the long-standing conjecture of [24] that the addition of the latter to the classic LP suffices for a constant integrality gap. Although this is not a result that concerns linear programming hierarchies, we included its presentation in that part of the thesis because the methodology we use is inspired by the local-to-global method and thus our proof deviates from standard integrality gap constructions. In fact we take the idea of fooling local constraints a little further: the bad solution fools every inequality π because its part that is *visible* to π, i.e., the variables in the support of π, can be extended to a solution that is a convex combination of feasible integer solutions for that instance or it is a convex combination of feasible solutions to another instance for which the same inequality is valid. Our proof relies on simple structural properties of the inequalities, disregarding the exact coefficients of the variables.

1.3 Extended Formulations: the currently most general mode of LP computation

In the past few years there has been an increasing interest in exposing the limitations of compact LP formulations for combinatorial optimization problems. The goal is to show a lower bound on the size of *extended formulations (EFs)* for a particular problem. Extended formulations add extra variables to the natural problem space; the increase in dimension may yield a smaller number of facets. The minimum size over all extended formulations is the *extension complexity* of the corresponding polytope. A superpolynomial lower bound on the extension complexity is of intrinsic interest in both polyhedral combinatorics and combinatorial optimization and implies that there is no polynomial-time algorithm relying purely on the solution of a compact linear program.

In the seminal paper of Yannakakis [35] the problem of lower bounding the size of extended formulations was considered for the first time: exponential lower bounds were proved for symmetric extended formulations of the matching and TSP polytopes. Yannakakis [35] identified also a crucial combinatorial parameter, the *nonnegative rank* of the slack matrix of the underlying polytope P, and he showed that it equals the extension complexity of P. A strong connection of the extension complexity of a polytope to communication complexity was made in [35], by showing that the nonnegative rank of the slack matrix is at least the size of its minimum rectangle cover. That connection has been exploited in several results on the extension complexity of polytopes.

Fiorini et al. [14] lifted the symmetry condition on the result of [35] regarding the TSP polytope, giving the first example of a polytope with exponential extension complexity and thus answering a long-standing open problem of [35]. Recently, Rothvoß [29] removed the symmetry condition for the matching polytope as well, answering the second long-standing open question of [35]. This was done by a breakthrough in bounding a refined version of the rectangle covering number.

A more general question is that of the size of approximate extended formulations. This problem was first considered in [10] where the methodology of [14] was extended to approximate formulations and an exponential bound for the linear encoding of the $n^{1/2-\varepsilon}$-approximate clique problem was given.

In [11] it was proved that in terms of approximating maximum constraint satisfaction problems (CSPs), LPs of size $O(n^k)$ are exactly as powerful as $O(k)$-level relaxations in the Sherali-Adams hierarchy. Their proof differs from previous work in showing that polynomials of low degree can approximate the functional version of the factorization theorem of [35].

Our contribution on Extended Formulations In the relevant part of the thesis we propose a new intuitive, geometric approach for proving lower bounds on the size of approximate extended formulations that relies on an insight on the expressive strength of “strong” sets of variables and encodings. Our contribution is summarized by the following.

First we introduce two very strong families of extended formulations (or relaxations) of a given polytope which we call *product formulations* and *distributional formulations*. The product relaxations are inspired by the study of the Sherali-Adams hierarchy – the variables have the intuitive meaning of corresponding to products over sets of variables from the original space. The distributional are intended to encode the problem in such a way that the feasible region is a straightforward distribution of (convex combination of) feasible integer solutions. (See Section 2 for the necessary definitions).

We prove in Theorem 1 that for any ρ-approximate extended formulation of a 0-1 polytope there is a product (distributional) formulation of the same size that is at least as strong. Theorem 1 reduces lower bounding the size of an extended formulation, which uses some unknown space and encoding, of a polytope P, to lower bounding the size of product (distributional) formulations of P. In the product (distributional) space we have the concrete advantage of knowing the section of the target relaxation. We extend the definition of product relaxations and our methodology to mixed integer sets. However in this case we are able to show that *mixed product relaxations* are at least as powerful as a special family of extended formulations.

Then we propose a methodology for proving lower bounds for relaxations for which the encoding of solutions is known, and in particular for product (distributional) formulations. The method is the following: first we define a set of vectors in the space of the relaxation such that for each one of those vectors there is an admissible objective function witnessing an integrality gap of ρ. We call that set of vectors the *core*. Then we show that, for any partition of the core into fewer than κ parts, there must be some part containing a set of conflicting vectors. A set of infeasible vectors is *conflicting* if its convex hull has nonempty intersection with the convex hull of $\{z^x \mid x \in P(x) \cap \{0,1\}^n\}$, which is always included in the feasible region of a product relaxation – here z^x is the encoding of feasible solution x to the variables of product formulations. Thus, we get that at least κ inequalities are needed to separate the members of the core from the feasible region and so κ is a lower bound on the size of any ρ-approximate product formulation. By considering the hypergraph whose set of vertices corresponds to the aforementioned set of vectors and whose set of hyperedges corresponds to the sets of conflicting vectors, the chromatic number of the hypergraph is a lower bound on the size of every ρ-approximate extended formulation. Moreover, there is always a core such that the chromatic number of the resulting, possibly infinite, hypergraph equals the extension complexity of the polytope at hand. Thus we give a characterization of extension complexity which can be seen as an alternative to the nonnegative rank of the slack matrix.

We exhibit a concrete application of our methodology by proving an exponential lower bound on the size of any $O(N)$-approximate mixed product relaxation for the CFL polytope, where N is the number of facilities in the instance. This result can be shown to imply that the $\Omega(N)$-level SA relaxation for CFL, which is obtained from any starting LP of size $2^{o(N)}$ defined on the classic set of variables, has unbounded gap $\Omega(N)$. This settles the open question of [4] whether there are LP relaxations upon which the application of lift-and-project methods captures the strength of preprocessing steps for CFL. This result establishes for the first time such a trade-off for a SA procedure that is independent of the starting relaxation K.

1.4 The strength of generalized configuration linear programs for capacitated versions of the facility location problem

In the relevant part of this thesis we introduce and study the family of proper relaxations which are configuration-like linear programs. The so-called *Configuration LP* was used by Bansal and Sviridenko [9] for the Santa Claus problem and has yielded valuable insights, mostly for resource allocation and scheduling problems (e.g., [31]). The analogue of the Configuration LP for facility location already exists, it is the *star relaxation* (see, e.g., [16]). In a star relaxation every variable corresponds to a *star,* i.e., a facility f and a set of clients assigned to f. The natural star relaxation for CFL and LBFL is equivalent to the standard LPs so it has an unbounded integrality gap. We generalize the idea of a star by introducing what we call *classes.* A *class* consists of a set with an arbitrary number of facilities and clients together with an assignment of each client to a facility in the set. The definition of a class can thus vary from simple, "local" assignments of clients to a single facility, to "global" snapshots of the instance that express the assignment of clients to a large set of facilities. A *proper relaxation* for an instance is defined by a collection $\mathcal{C}$ of classes and a decision variable for every class. We allow great freedom in defining $\mathcal{C}$; the only requirement is that the resulting formulation is symmetric and valid. The *complexity* α of a proper relaxation is the maximum fraction of the available facilities that are contained in a class of $\mathcal{C}$. Proper LPs are stronger than the standard relaxation. One can easily construct infinite families of instances where, by increasing the complexity in a proper relaxation, one cuts off more and more fractional solutions. We characterize the behavior of proper relaxations for CFL and LBFL through a sharp threshold result: anything less than maximum complexity results in a gap that is not bounded by any constant, while there are proper relaxations of maximum complexity with a gap of 1.

1.5 Publications

The publications that resulted from the work presented in this thesis include the following: [20] contains the results on the SA hierarchy, the flow-cover inequalities and the proper relaxations. The results regarding the LS hierarchy combined with the results of [19] and [20] were published in [21]. The results regarding the Extended Formulations are contained in [18].

2 Preliminaries on Extended Formulations

Given a polyhedron $K(x,y) = \{(x,y) \in \mathbb{R}^d \times \mathbb{R}^{d_y} \mid Ax + By \leq b\}$ the *projection to the x-space* is defined as $\{x \in \mathbb{R}^d \mid \exists y \in \mathbb{R}^{d_y} \colon Ax + By \leq b\}$, denoted as $\mathrm{proj}_{\mathrm{x}}(K(x,y))$. An *extended formulation* of a polyhedron $P(x) \subseteq \mathbb{R}^d$ is a linear system $K(x,y) = \{(x,y) \in \mathbb{R}^d \times \mathbb{R}^{d_y} \mid Ax + By \leq b\}$ such that $\mathrm{proj}_{\mathrm{x}}(K(x,y)) = P(x)$. The *size* of a polyhedron $P(x)$ is the minimum number of inequalities in its halfspace description. The *extension complexity* of $P(x)$ is the minimum size of an extended formulation of $P(x)$.

We define now ρ-approximate formulations as in [10]. Given a combinatorial optimization problem $T(S,f)$, a *linear encoding* of T is a pair (L,O) where $L \subseteq \{0,1\}^*$ is the set of encodings of *feasible solutions* to the problem and $O \subset \mathbb{R}^*$ is the set of encodings of the *admissible objective functions.* An instance of the linear encoding is a pair (d,w) where d is a positive integer defining the dimension of the instance and $w \subseteq O \cap \mathbb{R}^d$ is the set of admissible cost functions for instances of dimension d. Solving the instance (d,w) means finding $x \in L \cap \{0,1\}^d$ such that $w^T x$ is either maximum or minimum, according to the type of problem T. Let $P = \mathrm{conv}(\{x \in \{0,1\}^d \mid x \in L\})$ be the corresponding 0-1 polytope of dimension d. Given a linear encoding (L,O) of a maximization problem, the corresponding polytope P, and $\rho \geq 1$, a *ρ-approximate extended formulation of P* is an extended relaxation $Ax + By \leq b$ of P with $x \in \mathbb{R}^d, y \in \mathbb{R}^{d_y}$ such that

$$\begin{aligned}
\max\{w^T x \mid Ax + By \leq b\} &\geq \max\{w^T x \mid x \in P\} && \text{for all } w \in \mathbb{R}^d \text{ and} \\
\max\{w^T x \mid Ax + By \leq b\} &\leq \rho \max\{w^T x \mid x \in P\} && \text{for all } w \in O \cap \mathbb{R}^d.
\end{aligned}$$

For a minimization problem, we require

$$\min\{w^T x \mid Ax + By \le b\} \le \min\{w^T x \mid x \in P\} \qquad \text{for all } w \in \mathbb{R}^d \text{ and}$$
$$\min\{w^T x \mid Ax + By \le b\} \ge \rho^{-1} \min\{w^T x \mid x \in P\} \qquad \text{for all } w \in O \cap \mathbb{R}^d.$$

The *ρ-approximate extension complexity* of 0-1 integer polytope $P(x) \subseteq [0,1]^d$ is the minimum size of a ρ-approximate extended formulation of P.

We turn now to define a generic extended formulations that will play a central role.

Definition 1. *Given a* 0-1 *integer polytope $P(x) \subseteq [0,1]^d$, a* product formulation $D(z)$ of $P(x)$ *is an extended formulation $D(z)$ of $P(x)$, where $z \in \mathbb{R}^{2^d-1}$ and for every nonempty subset $\mathcal{E} \subseteq \{x_1, x_2, \ldots, x_d\}$ of the original variables, we have a variable $z_{\mathcal{E}}$, (where $z_{\{x_i\}}$ denotes x_i, $i = 1, \ldots, d$). For any feasible integer solution $x^s \in P(x) \cap \{0,1\}^d$ the vector z^s, whose components are defined as $z^s_{\mathcal{E}} = 1$ iff all variables in $\mathcal{E}$ have value* 1 *in x^s and $z^s_{\mathcal{E}} = 0$ otherwise, is feasible for any product formulation $D(z)$ of $P(x)$. We will refer to z^s as the* encoding *of the feasible integer solution x^s in the product variables.*

Note that the lifted polytope obtained from some specific linear relaxation of the 0-1 polytope $P(x)$, at any level of the SA hierarchy, after linearization and before projection to the original variables, is a (mixed) product relaxation.

3 The expressive power of product relaxations

In this section we show the following. For every 0-1 polytope $P(x)$ and every (approximate) extended formulation $Q(x,y) = \{(x,y) \in \mathbb{R}^{d_x} \times \mathbb{R}^{d_y} \mid Ax + By \le b\}$ of $P(x)$ there is a product formulation T_Q which has the size of $Q(x,y)$ and is at least as strong in terms of approximability. Similarly, we show that there is a distributional formulation R_Q of the same as $Q(x,y)$ and at least as strong.

A *substitution T for the product space* is a linear map of the form $y = Tz$ where T is a $d_y \times (2^{d_x} - 1)$ matrix and z is a $2^{d_x} - 1$ dimensional vector having a coordinate $z_{\mathcal{E}}$ for each nonempty set $\mathcal{E}$ of the form $\{x_i \mid i \in S \subseteq 2^{\{1,\ldots,d_x\}}\}$. For any substitution T, the *translation* of $Q(x,y)$, denoted T_Q, the formulation resulting by substituting $T_{(i)}z$, for y_i, $i = 1, \ldots, d_y$. Here $T_{(i)}$ denotes the ith row of T. We require that T_Q is a product formulation (see Definition 1) and we say that we have a *translation of Q to product formulations* (recall that the original variables x_i coincide with the variables $z_{\{x_i\}}$). Observe that the number of inequalities of T_Q is the same as in $Q(x,y)$. The translation may heighten exponentially the dimension, but, since our methodology will give lower bounds on the size of the product formulations, those bounds apply to the size of $Q(x,y)$ as well. A *substitution T for the distributional space* and a *translation to distributional formulations* is defined similarly.

Theorem 1. *Given a* 0-1 *polytope $P(x) \subseteq [0,1]^{d_x}$, for every polytope $Q(x,y)$ such that $P(x) \subseteq \mathrm{proj}_{\mathrm{x}}(Q(x,y))$ there is a translation T_Q to product formulations such that $P(x) \subseteq \mathrm{proj}_{\mathrm{x}}(T_Q) \subseteq \mathrm{proj}_{\mathrm{x}}(Q(x,y))$.*

Proof. We shall give a substitution T for the variables $y \in \mathbb{R}^{d_y}$ of $Q(x, y)$ so that the theorem holds. Let $g(x)$ be a section of $Q(x, y)$ (recall that a section associates every feasible 0-1 vector x of $P(x)$ to a specific y such that $(x, y) \in Q(x, y)$). We denote by $(p, 1) \in \mathbb{R}^{n+1}$ the vector resulting from $p \in \mathbb{R}^n$ by appending the scalar 1 as an extra coordinate.

Observe that a product variable $z_{\mathcal{E}}$ behaves, as far as the encondings z^s of solutions $x^s \in P(x) \cap \{0,1\}^{d_x}$ to product variables are concerned, like the monomial $\prod_{x_i \in \mathcal{E}} x_i$ would. Those monomials plus the constant 1 form the *Fourier basis*. Likewise we can see a variable y_i, as far as the encondings x^s, y^s of solutions $x^s \in P(x) \cap \{0,1\}^{d_x}$ are concerned, as a boolean function $y_i(x) : \{0,1\}^{d_x} \to \mathbb{R}$ such that $y_i(x^s) = y_i^s$. By basic functional analysis (see, e.g.,[17]), we have that every boolean function $y_i(x)$ has a unique Fourier representation $y_i(x) = \sum_{\mathcal{E} \subseteq \{x_i | i=1,\ldots,d_x\}} a_{\mathcal{E}}^{y_i} \prod_{x_i \in \mathcal{E}} x_i$. The intuition is that we will use the encodings z^s to product variables to simulate the encodings y^s. So we define the substitution T_i for a variable y_i as follows:

$$y_i = \sum_{\mathcal{E} \subseteq \{x_i | i=1,\ldots,d_x\}} a_{\mathcal{E}}^{y_i} z_{\mathcal{E}} \tag{1}$$

In the above expression we assume, for notational convenience that, $z_\emptyset = 1$. Recall that product variables are defined for nonempty sets.

Obviously $\mathrm{proj}_{\mathrm{x}}(T_Q) \subseteq \mathrm{proj}_{\mathrm{x}}(Q(x, y))$: from any feasible solution (x^0, z^0) of T_Q we can derive a feasible solution (x^0, y^0) of $Q(x, y)$ by setting y^0 equal to Tz^0.

We will now show that $P(x) \subseteq \mathrm{proj}_{\mathrm{x}}(T_Q)$ or, more specifically, that the encodings z^s of solutions to product variables are feasible for T_Q as required by the definition of product relaxations. Observe that by letting the z vector take the value z^s for some $s \in P \cap \{0,1\}^{d_x}$, by (1) we get that the quantities involved in the inequalities of T_Q are the exact same quantities involved in the corresponding inequalities of $Q(x, y)$ for $(x, y) = (x^s, y^s)$. By definition (x^s, y^s) is feasible for $Q(x, y)$ and thus z^s is feasible for T_Q.

Corollary 1. *A lower bound b on the size of any product relaxation D which is a ρ-approximate extended formulation of the* 0-1 *polytope $P(x)$, for $\rho \geq 1$, implies a lower bound b on the size of any ρ-approximate extended formulation $Q(x, y)$ of $P(x)$.*

4 Conclusion

In the context of this thesis we exposed the limitations of linear programming methods for providing satisfactory approximations to assignments problem with restrictions such as capacities. In particular we showed that the unboundedness of the integrality gap of CFL or LBFL relaxations persists even after applying the tightenings of the LS and SA hierarchies. We did so by proving the feasibility of a bad fractional solution for an asymptotically tight number of levels. We also proved that the submodular inequalities do not reduce the integrality gap to constant. Then, while turning our attention to the more general model of extended formulations, we devised a methodology for lower bounding the extension complexity which also serves as a characterization of the extension complexity. We applied our method to derive tight bounds on the size of mixed product relaxations which result also implies tight SA gaps regardless of the initial relaxation. Lastly, we proved similar negative results for families of proper relaxations that capture general configuration LPs. The obtained results answered a number of interesting open questions and conjectures from the relevant literature.

In the recent work of An et al. [5] the first constant factor LP-based approximation algorithm for CFL was given. However, the proposed relaxation is exponential in size and, according to the authors, it is not known to be separable in polynomial time. A natural question that arises is whether there is a polynomially-sized relaxation achieving a constant integrality gap. An interesting direction is that of determining the minimum size of an approximate extended formulation of the CFL polytope, which our results arguably suggest to be exponential. We leave this as an open problem.

Regarding our lower bounding methodology for extended formulations, the proof of our result for mixed product relaxations for CFL made use of a core whose underlying hypergraph is actually a graph and moreover a clique. To generalize this result to product formulations or distributional formulations, or to prove bounds on the extension complexity of other polytopes, we believe that the power of general hypergraphs needs to be exploited. Observe that our methodology requires only the existence of a suitable core, and thus, one could possibly employ probabilistic arguments to prove the existence of suitable hypergraphs of high chromatic number.

In the case of mixed integer polytopes, we leave as an open problem whether the mixed product relaxations are strong enough to simulate any extended formulation, as is the case for product relaxations and 0-1 polytopes.

We also believe that it would be interesting to revisit polytopes, whose extension complexity has been shown to be large, and provide independent proofs using our method, ideally by improving on the known bounds. Moreover, as we showed for CFL using product or distributional formulations one can provide lower bounds as well and this can be of help in settling the extension complexity.

References

1. Karen Aardal, Yves Pochet, and Laurence A. Wolsey. Capacitated facility location: Valid inequalities and facets. *Mathematics of Operations Research*, 20:562–582, 1995.
2. Ankit Aggarwal, Anand Louis, Manisha Bansal, Naveen Garg, Neelima Gupta, Shubham Gupta, and Surabhi Jain. A 3-approximation algorithm for the facility location problem with uniform capacities. *Mathematical Programming*, 141(1-2):527–547, 2013.
3. Sara Ahmadian and Chaitanya Swamy. Improved approximation guarantees for lower-bounded facility location. *CoRR*, abs/1104.3128, 2011.

4. Hyung-Chan An, Aditya Bhaskara, and Ola Svensson. Centrality of trees for capacitated k-center. *CoRR*, abs/1304.2983, 2013.
5. Hyung-Chan An, Mohit Singh, and Ola Svensson. LP-based algorithms for capacitated facility location. In *55th IEEE Annual Symposium on Foundations of Computer Science, FOCS 2014, Philadelphia, PA, USA, October 18-21, 2014*, pages 256–265. IEEE Computer Society, 2014.
6. Sanjeev Arora, Béla Bollobás, László Lovász, and Iannis Tourlakis. Proving integrality gaps without knowing the linear program. *Theory of Computing*, 2(1):19–51, 2006.
7. Egon Balas, Sebastián Ceria, and Gérard Cornuéjols. A lift-and-project cutting plane algorithm for mixed 0-1 programs. *Mathematical Programming*, 58(3):295–324, February 1993.
8. Manisha Bansal, Naveen Garg, and Neelima Gupta. A 5-approximation for capacitated facility location. In Leah Epstein and Paolo Ferragina, editors, *Algorithms ESA 2012*, volume 7501 of *Lecture Notes in Computer Science*, pages 133–144. Springer Berlin Heidelberg, 2012.
9. Nikhil Bansal and Maxim Sviridenko. The Santa Claus problem. In *Proceedings of the thirty-eighth annual ACM symposium on Theory of computing*, STOC '06, pages 31–40, New York, NY, USA, 2006. ACM.
10. Gábor Braun, Samuel Fiorini, Sebastian Pokutta, and David Steurer. Approximation limits of linear programs (beyond hierarchies). In *FOCS*, pages 480–489, 2012.
11. Siu On Chan, James R. Lee, Prasad Raghavendra, and David Steurer. Approximate constraint satisfaction requires large LP relaxations. In *54th Annual IEEE Symposium on Foundations of Computer Science, FOCS 2013, 26-29 October, 2013, Berkeley, CA, USA*, pages 350–359, 2013.
12. Gérard Cornuéjols. Valid inequalities for mixed integer linear programs. *Mathematical Programming*, 112(1):3–44, 2008.
13. Wenceslas Fernandez de la Vega and Claire Kenyon-Mathieu. Linear programming relaxations of maxcut. In *Proceedings of the eighteenth annual ACM-SIAM symposium on Discrete algorithms*, SODA '07, pages 53–61, Philadelphia, PA, USA, 2007. Society for Industrial and Applied Mathematics.
14. Samuel Fiorini, Serge Massar, Sebastian Pokutta, Hans Raj Tiwary, and Ronald de Wolf. Linear vs. semidefinite extended formulations: exponential separation and strong lower bounds. In *Proceedings of the 44th Symposium on Theory of Computing*, STOC '12, pages 95–106, New York, NY, USA, 2012. ACM.
15. Dorit S. Hochbaum. Heuristics for the fixed cost median problem. *Mathematical Programming*, 22:148–162, 1982.
16. Kamal Jain, Mohammad Mahdian, Evangelos Markakis, Amin Saberi, and Vijay V. Vazirani. Greedy facility location algorithms analyzed using dual fitting with factor-revealing LP. *Journal of the ACM*, 50(6):795–824, November 2003.
17. Stasys Jukna. *Boolean Function Complexity - Advances and Frontiers*, volume 27 of *Algorithms and combinatorics*. Springer, 2012.
18. Stavros G. Kolliopoulos and Yannis Moysoglou. Extended formulation lower bounds via hypergraph coloring? *to appear in 32nd International Symposium on Theoretical Aspects of Computer Science, STACS 2015, Munich, Germany, March 4th -7th*, 2015.
19. Stavros G. Kolliopoulos and Yannis Moysoglou. Tight bounds on the Lovász-Schrijver rank for approximate capacitated facility location. *Manuscript*, 2013.

20. Stavros G. Kolliopoulos and Yannis Moysoglou. Sherali-Adams gaps, flow-cover inequalities and generalized configurations for capacity-constrained facility location. In *Approximation, Randomization, and Combinatorial Optimization. Algorithms and Techniques, APPROX/RANDOM 2014, September 4-6, 2014, Barcelona, Spain*, pages 297–312, 2014.
21. Stavros G. Kolliopoulos and Yannis Moysoglou. Integrality gaps for strengthened linear relaxations of capacitated facility location. *To appear in Mathematical Programming, 43 pages,*, 150(1), 2015.
22. Jean B. Lasserre. An explicit exact SDP relaxation for nonlinear 0-1 programs. In *Proceedings of the 8th International IPCO Conference on Integer Programming and Combinatorial Optimization*, pages 293–303, London, UK, UK, 2001. Springer-Verlag.
23. Monique Laurent. A Comparison of the Sherali-Adams, Lovász-Schrijver, and Lasserre relaxations for 0–1 Programming. *Mathematics of Operations Research*, 28(3):470–496, July 2003.
24. Retsef Levi, David B. Shmoys, and Chaitanya Swamy. LP-based approximation algorithms for capacitated facility location. *Mathematical Programming*, 131(1-2):365–379, 2012. Preliminary version in Proc. IPCO 2004.
25. Shi Li. A 1.488 approximation algorithm for the uncapacitated facility location problem. *Information and Computation*, 222:45–58, January 2013.
26. Shi Li and Ola Svensson. Approximating k-median via pseudo-approximation. In Dan Boneh, Tim Roughgarden, and Joan Feigenbaum, editors, *Proc. 45th STOC*, pages 901–910. ACM, 2013.
27. Lazlo Lovász and Alexander Schrijver. Cones of matrices and set-functions and 0-1 optimization. *SIAM Journal on Optimization*, 1:166–190, 1991.
28. Ran Raz and Shmuel Safra. A sub-constant error-probability low-degree test, and a sub-constant error-probability PCP characterization of NP. In *Proceedings of the 29th Annual ACM Symposium on Theory of Computing*, pages 475–484, 1997.
29. Thomas Rothvoß. The matching polytope has exponential extension complexity. In *Proc. STOC*, pages 263–272, 2014.
30. Hanif D. Sherali and Warren P. Adams. A hierarchy of relaxations between the continuous and convex hull representations for zero-one programming problems. *SIAM Journal of Discrete Mathematics*, 3(3):411–430, 1990.
31. Ola Svensson. Santa Claus schedules jobs on unrelated machines. In *Proceedings of the 43rd annual ACM symposium on Theory of computing*, STOC '11, pages 617–626, New York, NY, USA, 2011. ACM.
32. Z. Svitkina. Lower-bounded facility location. In *Proceedings of the 19th ACM-SIAM Symposium on Discrete Algorithms*, pages 1154–1163, 2008.
33. Jens Vygen. Approximation algorithms for facility location problems (Lecture Notes). Report 05950-OR, Research Institute for Discrete Mathematics, University of Bonn, 2005. URL: www.or.uni-bonn.de/∼vygen/files/fl.pdf.
34. David P. Williamson and David B. Shmoys. *The Design of Approximation Algorithms*. Cambridge University Press, New York, NY, USA, 1st edition, 2011.
35. Mihalis Yannakakis. Expressing combinatorial optimization problems by linear programs. *Journal of Computer and System Sciences*, 43(3):441 – 466, 1991.