DIMACS Technical Report 2004-56
December 2004

Controlled Linear Programming: Duality and Boundedness

by

Henrik Björklund, Olle Nilsson, Ola Svensson, Sergei Vorobyov [1] ,[2] ,[3] ,[4] ,[5]

[1]DIMACS visitors, November, 2004.
[2]Uppsala University, Information Technology Department, Box 337, 751 05 Uppsala, Sweden.
[3]Supported by the Grant IG2003-2 067 from the Swedish Foundation for International Cooperation in Research and Higher Education (STINT, www.stint.se) and the Swedish Research Council grants "*Infinite Games: Algorithms and Complexity*" and "*Interior-Point Methods for Infinite Games*".
[4]DIMACS hosts: Leonid Khachiyan and Endre Boros
[5]Email: vorobyov@csd.uu.se

ABSTRACT

We further investigate the Controlled Linear Programming Problem (CLPP) introduced in [DIMACS-TR-2004-41], by defining and studying the Dual CLPP, boundedness, duality, stability, optimality conditions, correctness of subexponential algorithms, conditions for NP∩coNP-membership.

Contents

1 Introduction

A *controlled optimization problem* is a succinctly represented family $\mathcal{F} = \{G_i\}_{i \in I}$ of instances of a polynomial time solvable optimization problem, with cost function c. A *controller* wants to find an instance $G_i \in \mathcal{F}$ with the largest value $opt(c, G_i)$, over all instances in $\mathcal{F}$.

In [4] we introduced the *Longest-Shortest Paths* controlled optimization problem. Let $G = (V, E)$ be a weighted digraph with a distinguished sink t and a subset $U \subseteq V$ of *controlled vertices*. In every such vertex, the controller is allowed to throw away all but one outgoing edges, with the goal to make the shortest paths to t as long as possible. Thus the set of subgraphs of G obtained by selecting exactly one outgoing edge from each controlled vertex, with the sum of shortest paths lengths as the associated cost function, is a controlled optimization problem. Investigating this problem allowed us to suggest the first combinatorial subexponential algorithms for mean payoff games [4].

The *Controlled Linear Programming Problem (CLPP)* is another example, first introduced and investigated in [1]. A CLPP instance consists of a number of linear inequalities of certain form. The *controller* has the opportunity to select and discard constraints according to simple rules, with the goal to make the optimal value of the cost function $\sum_{v \in Var} v$ in the remaining linear program as large as possible.

The CLPP, like the longest-shortest paths problem, is closely related to a number of infinite games on finite graphs, such as parity, mean payoff, discounted payoff, and simple stochastic games, all of considerable importance in the fields of computer-aided verification and algorithmic game theory. In fact, all of these games can be adequately modeled as CLPPs [1]. This makes the CLPP worthy of study for number of reasons. First, it gives us a simple unified view of strategy improvement algorithms for infinite games. Second, as the CLPP is described in terms of linear algebra and linear programming, the multitude of results from these fields can be used for investigating game problems, a work we began in [1] and continue here. Third, the CLPP problem leads to the interesting question of how much the games mentioned above can be generalized while staying in NP∩coNP. This question is also given partial answers here and in [1].

The decision problems for parity, mean payoff, discounted payoff, and simple stochastic games all belong to NP∩coNP. There is a number of deterministic algorithms for solving them; see, e.g., [6, 5, 9, 12, 7, 11], but none has been shown to have subexponential complexity. There are, however, randomized algorithms which solve the problems in expected subexponential time, as shown in [8] for binary simple stochastic games, and in [2, 3, 4, 1] for parity, mean payoff, discounted payoff, and general outdegree simple stochastic games.

The CLPP is more general than the games discussed above, and it is not known whether it can be optimized, in general, in expected subexponential time or even belongs to NP∩coNP. In [1], however, we show that this is the case for a number of possible extensions of the games. This work is set forth in the present paper. We investigate which conditions on the structure of the constraints are sufficient for subexponential solvability and membership in NP∩coNP. To this end, we also define a dual minimization version of the CLPP, and study how the relation between an instance and its dual can be used to infer boundedness of an optimal solution.

The crucial property underlying the subexponential iterative optimization for CLPP is "stability $\Rightarrow$ optimality". A strategy (choice of constraints), which cannot be improved by local changes, is called *stable* (stability is polynomial time verifiable). Whenever stable strategies are guaranteed to be optimal, our subexponential algorithms are applicable and correct. A number of CLPP classes satisfying "stability $\Rightarrow$ optimality", covering all games mentioned above, are investigated in [1].

In this paper, we derive the desired "stability $\Rightarrow$ optimality" property using the *duality* and *boundedness* considerations. This leads to conceptually easier and more transparent proofs, compared to the proofs from [1], based on potential transformations of hypergraphs and convergence of infinite series. We also think that introducing the concept of a Dual CLPP and investigating it together with the Primal makes the theory more symmetric and leads to a better understanding of the problem. As a result, we close certain gaps in the previous development [1], like correctness of attractive switches, existence of optimal counterstrategies. We also find simple yet fundamental relations between linear algebraic concepts and games, including nonsingularity, Farkas lemma, etc.

2 The Controlled Linear Programming Problem

In this section we present the controlled linear programming problem (CLPP), first defined in [1], and define the dual minimization problem. We also review some fundamental properties of CLPP.

2.1 Primal CLPP

Let $x = (x_1, \dots, x_n, x_{n+1}, \dots, x_{n+p})$ be a vector of variables. The first n variables are called *controlled*, and the remaining ones are *uncontrolled*. By p_i^j we denote linear *homogeneous* polynomials with *nonnegative* coefficients and variables in x, and by w_i^j real numbers. A *monotone* constraint has form $x_i \leq p_i^j + w_i^j$. Such a constraint is called:

1. *optional*, if x_i is a controlled variable, and
2. *compulsory*, if x_i is an uncontrolled variable.

Let S be a system of monotone constraints with at least one constraint per variable x_i in x. A *controller strategy* σ consists in selecting *exactly one* constraint for each controlled variable. Denote by $S(\sigma)$ the resulting system of constraints obtained by deleting from S all optional constraints not selected by σ.

Definition 2.1 *An instance of the* Controlled Linear Programming Problem *(CLPP for short) is specified as follows.*

Given: *a system of monotone constraints S.*

Find: *a controller strategy σ^* that* maximizes $\max(1^T x | S(\sigma^*))$, *over all controller strategies.* □

2.2 Dual CLPP

Call a constraint of the form $x_i \geq p_i^j + w_i^j$ *anti-monotone*. As before, such a constraint is *optional* if x_i is controlled and *compulsory* otherwise. Let S^d be a system of anti-monotone constraints with at least one constraint per variable x_i in x.

Symmetrically, an *environment strategy* τ is a selection of exactly one constraint for each *uncontrolled* variable. Denote by $S^d(\tau)$ the resulting system of constraints obtained by deleting from S^d all compulsory constraints not selected by τ.

Definition 2.2 *An instance of the* Dual Controlled Linear Programming Problem *(DCLPP for short) is specified as follows.*

Given: *a system of anti-monotone constraints S^d.*

Find: *an environment strategy τ^* that* minimizes $\min(1^T x | S^d(\tau^*))$, *over all environment strategies.* □

In the sequel the controller and the environment will more often be referred to as players MAX and MIN, respectively. Given a constraint system S we will denote its dual by S^d. If both MAX and MIN fix their strategies σ and τ then the resulting systems will be denoted by $S(\sigma,\tau)$ and $S^d(\sigma,\tau)$. For a system $S(\sigma,\tau)$ or $S^d(\sigma,\tau)$ denote $S^=(\sigma,\tau)$ the system obtained by replacing all inequalities by *equalities*.

Notational Convention. To denote the optimal values of the linear program associated with a constraint system, we use the following notation.

$$\begin{aligned} opt(S(\sigma,\tau)) &=_{df} \max(1^T x | S(\sigma,\tau)), \\ opt(S(\sigma)) &=_{df} \max(1^T x | S(\sigma,\tau)), \\ opt(S^d(\sigma,\tau)) &=_{df} \min(1^T x | S^d(\sigma,\tau)), \\ opt(S^d(\tau)) &=_{df} \min(1^T x | S^d(\tau)). \end{aligned}$$

2.3 Admissibility of Strategies

A strategy σ of MAX is called *admissible* if $S(\sigma)$ is feasible. (It follows that for each strategy τ of MIN the system $S(\sigma,\tau)$ is feasible). Symmetrically, a strategy τ of MIN is admissible if $S^d(\tau)$ is feasible.

Notational Convention. Throughout the paper we systematically use Greek letters σ, τ for MAX's and MIN's strategies, respectively. □

2.4 Monotonicity of Iterative Improvement. Attractiveness and Profitability

For a strategy σ, a *single switch* changes the selected constraint for exactly one variable. Suppose the current strategy is σ, and variable assignment $x^* = \arg\max(1^T x | S(\sigma))$. Let $\sigma_i \equiv x_i \leq p_i^\sigma(x) + w_i^\sigma$ be the constraint selected for the controlled variable x_i in σ, and let $\sigma'_i \equiv x_i \leq p'_i(x) + w'_i$ be another constraint in S. The single switch from σ_i to σ'_i is called *attractive* if $p_i^\sigma(x^*) + w_i^\sigma < p'_i(x^*) + w'_i$.[1] A strategy without attractive switches is called *stable*. Attractive switches and stable strategies for MIN in the DCLPP are defined symmetrically.

In [1] we demonstrated that in any CLPP instance, every combination of attractive single switches for MAX is profitable, i.e., leads to a monotonic improvement of the target function. Here we prove a dual claim.

Theorem 2.3 *Any combination of attractive switches for* MIN *in the DCLPP instance system S^d is profitable, i.e., leads to a monotonic decrease of the target function $1^T x$.*

[1]This definition implicitly assumes that x^* is the only optimal solution to $S(\sigma)$. Otherwise, attractiveness of a switch depends on a selected optimal solution. In fact, by Corollary 4.2, the system only has one solution.

Proof. Assume that a number of attractive switches are made, from strategy τ to τ'. Let the variable assignment $x^* = \arg\max(1^T x | S^d(\tau))$. For every switched variable x_i, we have the old constraint $x_i \geq p_i^\tau(x) + w_i^\tau$, and the new constraint $x_i \geq p_i^{\tau'}(x) + w_i^{\tau'})$.

Because the switches are attractive, it is easy to see that x^* is a solution to the new LP $S^d(\tau')$, because every constraint is either the same or more generous, i.e., with a smaller right-hand-side.

Take one switched variable x_i, let $x'_i = x^*_i - \varepsilon$, for a small enough $\varepsilon > 0$, and let $x'_j = x^*_j$ for all $j \neq i$. Then x' is a solution to the new LP $S^d(\tau')$, since

1. x_i does not violate its new constraint, which by attractiveness is more generous, and
2. the decrease of x_i does not increase the right-hand-side polynomials of any constraint, since all coefficients are nonnegative.

We have thus shown that the resulting strategy is admissible and that there exists a better solution to the new LP. Therefore, the switch is profitable. In addition, we see that the values of variables where attractive switches are made are guaranteed to decrease in an optimal solution to the new system.

Note that the argument above essentially uses the fact that the right-hand-side polynomials of constraints have nonnegative coefficients. The argument fails if negative coefficients are allowed. □

2.5 Stability and Optimality

Recall that a strategy is *stable* if it has no attractive switches. We showed in [1] that for each CLPP instance S such that every stable strategy σ is also optimal, the value $\max_\sigma \max(1^T x | S(\sigma))$ can be found in *randomized subexponential* time. We also showed that this condition holds for a number of interesting subclasses, including

1. the *univariate, unit coefficient* case, with constraints on the form $x_i \leq x_i^j + w_i^j$, subsuming *mean payoff games* and, by reduction, *parity games*, and
2. the *discounted* case, with constraints $x_i \leq \sum_{j=1}^{n+p} c_i^j x_i^j + w_i^j$, such that $\sum_{j=1}^{n+p} c_i < 1$, subsuming *discounted payoff and stopping simple stochastic games.*

The subexponential analysis of the algorithms, adapted from combinatorial linear programming in [1], relies on the fact that once a local optimum of a subsystem (with a subset of the original constraints) has been found, we never need to return to a strategy in the subsystem. When stability does not imply optimality, this property no longer holds, and the algorithms might even fail to correctly find an optimum.

3 Fundamental Properties of CLPPs

In this section we prove a number of basic lemmas that we rely heavily upon in the sequel.

3.1 Tightness of Finite Optimal Solutions

Lemma 3.1 *Let S be a CLPP instance and σ, τ be a pair of strategies such that* $x^* = \arg\max(1^T x|S(\sigma,\tau))$ *is finite. Then*

1. *x^* is a feasible solution to $S^=(\sigma,\tau)$,*
2. *conversely, $1^T x' = 1^T x^*$ for every other feasible solution x' to $S^=(\sigma,\tau)$.*

   In other words, optimal solutions to $\max(1^T x|S(\sigma,\tau))$ *are tight, and tight solutions are optimal.*

Proof. If x^*, a feasible solution to $S(\sigma,\tau)$, did not solve the system $S^=(\sigma,\tau)$, it wouldn't be an optimal solution to $S(\sigma,\tau)$, since for some constraint, the inequality would be non-tight, and we could slightly increase the value of the corresponding variable, without violating any constraints (due to monotonicity). Suppose that $S^=(\sigma,\tau)$ has a solution x' such that $1^T x' \neq 1^T x^*$. Rewrite $S^=(\sigma,\tau)$ in matrix form as $Ax = b$. Since $Ax' = Ax^*$ we get $A(x' - x^*) = 0$. Let λ be any real number. Then $Ax^* + A\lambda(x' - x^*) = A(x^* + \lambda(x' - x^*)) = b$. Since $1^T x' \neq 1^T x^*$ this means that $Ax = b$ has arbitrarily large solutions, contradicting the boundedness assumption. □

3.2 Existence of an Optimal Response Counterstrategy

Lemma 3.2 *Let S be a CLPP and σ an admissible strategy in S. Then there is a counter-strategy τ of* Min *such that $opt(S(\sigma,\tau)) = opt(S(\sigma))$.*

Proof. If $opt(S(\sigma)) = +\infty$, the claim is trivial, because $S(\sigma,\tau)$ has less constraints.

Otherwise, rewrite the linear program $\max(1^T x|S(\sigma))$ as $\max(1^T x|Ax \leq b)$. Since its optimal value M is finite, for each solution satisfying $Ax \leq b$, the inequality $1^T x \leq M$ holds. By the affine form of Farkas Lemma (see, e.g., [10, Corollary 7.1h]), we have that $1^T x \leq M$ is a positive linear combination of some constraints in $Ax \leq b$.

Let n be the number of variables in $S(\sigma)$. Then there exists at most n linearly independent constraints, and since the solution is bounded there must exist at least one constraint per variable. Thus we are able to choose one constraint per variable that are linearly independent such that $1^T x \leq M$ is a positive linear combination of the chosen constraints. The new linear program $max(1^T x|A'x \leq b')$ has the optimal value M, because $1^T x \leq M$ and $A'x \leq b'$ is feasible. Now take the constraint selected for every compulsory variable as the strategy τ and $M = opt(S(\sigma,\tau))$. □

4 Boundedness of Arbitrary CLPPs

In this section we investigate necessary and sufficient conditions for boundedness. Interestingly, boundedness is related to well-known linear algebraic concepts, like nonsingularity and kernels of linear operators.

4.1 Nonsingularity: A Necessary Condition of Boundedness

It turns out that the player Min, in order to be able to keep the optimal value of the CLPP instance *finite*, has to select counterstrategies in such a way that the resulting constraint matrix is *nonsingular*.

Lemma 4.1 *Suppose σ is admissible in S and $opt(S(\sigma)) < +\infty$. Then for every strategy τ of* Min *for which $opt(S(\sigma, \tau)) < +\infty$, the matrix A of the constraint system $S(\sigma, \tau)$, written as $Ax \leq b$, is nonsingular.*

Proof. Let $x^* = \arg\max(1^T x | Ax \leq b)$. Such a vector exists by admissibility of σ.

Assume, toward a contradiction, that A is singular. Then there is a nontrivial solution $y^* \neq 0$ to $Ax = 0$. Suppose $1^T y^* \neq 0$. Then $Ax^* + \lambda A y^* = A(x^* + \lambda y^*) \leq b$ for any $\lambda \in \mathbb{R}$. This contradicts the assumption that x^* is optimal. Thus $1^T y^* = 0$.

Let m be the number of variables in x, and let $I \subseteq \{1, \ldots, m\}$ be the set of indices such that $y_i^* \leq 0$ for all $i \in I$ and $y_j^* > 0$ for all $j \notin I$. Notice that I is a nontrivial subset of $\{1, \ldots, m\}$. Now construct y' in the following way. If $i \in I$, then set $y'_i = 0$, otherwise $y'_i = y_i^*$. Consider a constraint from $Ax \leq 0$. It has the form $x_i \leq p(x)$, where p is a homogeneous polynomial with nonnegative coefficients. We know that $Ay^* \leq 0$. If $i \notin I$, then $y'_i = y_i^*$ and $p(y') \geq p(y^*)$, so y' satisfies the constraint. On the other hand, if $i \in I$, then $y'_i = 0$, and since all entries in y' are nonnegative, $p(y') \geq 0$, so the constraint is satisfied. Thus $Ay' \leq 0$.

This means that $Ax^* + Ay' = A(x^* + y') \leq b$ and $1^T y' > 0$, thus $1^T(x^* + y') > 1^T x^*$, which contradicts the assumption that x^* is optimal. Thus 0 is the unique solution to $Ax = 0$ and A is nonsingular. □

Remark. Note that the special structure of CLPP (positiveness of constraints) is essentially used in the proof above. Indeed, consider the matrix

$$A = \begin{pmatrix} 1 & 1 \\ 1 & 1 \end{pmatrix},$$

which *does not correspond to any CLPP instance*. Although $\max(1^T x | Ax \leq 0) = 0$, the matrix A is singular. For arbitrary such (non-CLPP-generated) matrices we can only claim that the kernel of the associated linear operator is a linear subspace of the hyperplane $\{x | 1^T x = 0\}$, once the optimum is finite.

4.2 Bounded Systems Have Unique Solutions

In Section 2.4 we defined attractive switches for a strategy σ as those that seem to give an improvement with respect to the value vector $x^* = \arg\max(1^T x | S(\sigma))$. In the definition the optimal solution was implicitly assumed to be unique. We are now ready to justify this assumption.

Corollary 4.2 *Suppose σ is admissible in S and opt$(S(\sigma)) < +\infty$. Then* $\arg\max(1^T x|S(\sigma))$ *is* uniquely *defined.*

Proof. Let τ be an optimal counterstrategy to σ, and rewrite $S(\sigma, \tau)$ as $Ax \leq b$. If x^* is an optimal solution to $Ax \leq b$, then we must have $Ax^* = b$ by Lemma 3.1, and since A is nonsingular by Lemma 4.1, this solution is unique. □

4.3 Nonsingularity Does not Imply Boundedness

An attempt to invert Lemma 4.1 fails, as is shown in the following

Proposition 4.3 *Nonsingularity of the constraint matrix (resulting from fixed player strategies) does not imply boundedness.*

Proof. Let strategies of Max and Min result in the constraints

$$\begin{array}{rcl} x_1 & \leq & x_2 \\ x_2 & \leq & 2 \cdot x_1 \end{array}$$

The matrix of this system is nonsingular:

$$A = \begin{pmatrix} 1 & -1 \\ -2 & 1 \end{pmatrix} \text{ and } A^{-1} = \begin{pmatrix} -1 & -1 \\ -2 & -1 \end{pmatrix},$$

and the unique tight solution to the system is $x_1, x_2 = 0$. But $Ax \leq b$ has arbitrarily large solutions, e.g., $(1, 1)$, $(2, 2)$, $(3, 3)$, ..., none of which satisfies both constraints as equalities. □

4.4 Sufficient Condition for Boundedness

An additional requirement to nonsingularity allows to conclude boundedness.

Lemma 4.4 *Suppose for every pair of strategies σ, τ of Max, Min the matrix A for the constraint system $S(\sigma, \tau)$, written in the form $Ax \leq b$, satisfies:*

1. *A is nonsingular;*

2. *in the inverse matrix A^{-1} every column sums up to a nonnegative number.*

Then $\max_\sigma \ opt(S(\sigma)) < +\infty$.

Proof. Consider the set of all solutions to the system $Ax \leq b$, which may be equivalently rewritten as $Ax + y = b$ with $y \geq 0$ for fresh variables y (i.e., y is a nonnegative vector of slack variables). Rewriting it equivalently as

$$x = -A^{-1}y + A^{-1}b, \tag{1}$$

we see that for every nonnegative vector $y^* \geq 0$ the solution x^* to (1) is *uniquely* determined. Now the question reduces to whether

$$\max(1^T x | x = -A^{-1}y, \ y \geq 0) < +\infty \tag{2}$$

(Note that we drop out the constant vector $A^{-1}b$ as nonsignificant for unboundedness.)

The question (2) is easy to answer. $-A^{-1}$ is a nonsingular linear operator transforming a nonnegative orthant (cone) $y \geq 0$ to another unbounded full dimensional cone with 0 as origin. The question (2) therefore becomes:

> whether following any of the rays $-A^{-1}\lambda e_i$, for $\lambda > 0$ and e_i the i-th basis vector (1 in i-th position, zeros elsewhere) allows to unboundedly increase the sum of all coordinates?

It holds iff there is a column i in A^{-1} such that the sum of elements in this column is negative, and the claim follows. □

5 Strong Boundedness, Uniqueness, and Duality

In this section we introduce the concept of a strongly bounded CLPP, by stipulating finiteness of target values in primal and dual instances, for any choice of admissible strategies by the players. We show that strong boundedness implies equality of optimal values of the primal and dual CLPPs. As corollaries, we derive the desirable properties that every stable strategy is optimal, and the subexponential algorithms are valid for both primal and dual problems.

5.1 Strong Boundedness

Definition 5.1 (Strong Boundedness) *Say that a CLPP instance S is* strongly bounded *if for every admissible strategies σ of Max in S and τ of Min in S^d one has:*

$$\begin{aligned} opt(S(\sigma,\tau)) &< +\infty, \\ opt(S^d(\sigma,\tau)) &> -\infty. \end{aligned}$$

As a direct consequence of the definition and Lemma 4.1, we have the following

Proposition 5.2 *For a strongly bounded CLPP instance, and any choice of admissible strategies by the players, the resulting constraint matrix is nonsingular.*

Although strong boundedness may seem unrealistically restrictive, it nevertheless covers and immediately applies to stopping simple stochastic and discounted payoff games, since strong boundedness automatically holds for such games. Note that strong boundedness does not hold for CLPP instances acquired from mean payoff games.

5.2 Minimax Principle

We start by a useful general property of strongly bounded CLPPs, equating the optimal values in the primal and dual CLPPS, for any fixed admissible strategies of the players.

Proposition 5.3 (Minimax Principle) *For every strongly bounded system S and any admissible strategies σ of Max in S and τ of Min in S^d one has $opt(S(\sigma,\tau)) = opt(S^d(\sigma,\tau))$.*

Proof. Assume $opt(S(\sigma,\tau)) = M \neq N = opt(S^d(\sigma,\tau))$. Since $M < +\infty$, by Lemma 3.1, an optimal solution u^* witnessing for M satisfies $S(\sigma,\tau)$ as equalities $S^=(\sigma,\tau)$, and $M = \sum_i u_i^*$. Similarly, v^* witnessing N satisfies $S^d(\sigma,\tau)$ as equalities $S^=(\sigma,\tau)$, and $N = \sum_i v_i^*$. Since u^* is different from v^*, $u^* + \alpha(u^* - v^*)$ is a solution to the system $S^=(\sigma,\tau)$ for *every* $\alpha \in \mathbb{R}$. Indeed, writing $S^=(\sigma,\tau)$ in the form $Ax = b$, it is easy to verify that $A(u^* + \alpha(u^* - v^*)) = b$. Note that $\sum_i (u_i^* - v_i^*) = M - N \neq 0$ by assumption.

Selecting a negative or positive α big enough we can make M arbitrary big (pick any $\alpha > 0$ if the sum of coordinates in $u^* - v^*$ is positive and any $\alpha < 0$ otherwise). This contradicts to the boundedness of M. Therefore, $M = N$. □

5.3 Stability and Optimality

Lemma 5.4 *Let σ be stable in a strongly bounded S, and τ be an optimal response to σ (see Lemma 3.2). Then*

$$opt(S(\sigma)) = \max_{\sigma} \; opt(S(\sigma,\tau)).$$

In other words, Max cannot improve the value even if Min invariably uses τ.

Proof. The $\leq$ relation is obvious.

Suppose, Min keeps his strategy τ, and in σ (that selects tight constraints) Max makes a *flat* multiple switch (resulting in a new strategy σ') to other constraints satisfied as equalities in an optimal solution x^*. Lemma 5.5 demonstrates that a flat switch does not change the optimum value.

Assume, Min keeps his strategy τ, but Max makes an *anti-attractive switch.* Lemma 5.6 shows that such a switch does not improve the optimum value.

Since an arbitrary nonattractive switch is a combination of a flat and an anti-attractive switch, the conclusion follows. □

Lemma 5.5 *Assuming strong boundedness of S, every flat switch for Max in $S(\sigma,\tau)$ (only Max has choices) preserves the optimal value. Symmetrically, every flat switch for Min in $S^d(\sigma,\tau)$ (only Min has choices) preserves the optimal value.*

Proof. By Lemma 3.1, $u^* = \arg\max(1^T x | S(\sigma,\tau))$ satisfies $S^=(\sigma,\tau)$. After making a flat switch, to another strategy σ', the vector u^* is still a solution to $S^=(\sigma',\tau)$. Again by Lemma 3.1[2], u^* is an optimal solution to $\max(1^T x | S(\sigma',\tau))$. □

Lemma 5.6 *An anti-attractive switch for Max in $S(\sigma,\tau)$ (only Max has choices) does not increase the optimal value. Symmetrically, an anti-attractive switch for Min in $S^d(\sigma,\tau)$ (only Min has choices) does not decrease the optimal value.*

Proof. By Lemma 3.1, $u^* = \arg\max(1^T x | S(\sigma,\tau))$ satisfies $S^=(\sigma,\tau)$. After making an anti-attractive switch, to another strategy σ', the vector u^* is still a *feasible* (but non-tight) solution to the dual system $S^d(\sigma',\tau)$, because such a switch corresponds to the following system of inequalities:

$$\begin{array}{rcl} x & = & p \text{ (in } \sigma\text{)}, \\ & & > \\ x & > & p' \text{ (in } \sigma'\text{)}. \end{array}$$

Therefore, the optimal solution to $S^d(\sigma',\tau)$ is smaller (in value) and the claim follows by Proposition 5.3. □

Corollary 5.7 *Every stable strategy in a strongly bounded CLPP instance is optimal.*

[2]Here we essentially rely on strong boundedness. In general a switch to σ' may lead to unboundedness.

Proof. Suppose we have a stable strategy, but there is a better one. It can be obtained by a multiple switch. This switch decomposes in two switches, one flat, and one anti-attractive. By the above lemmas, the first one preserves the value and the second one decreases it. The claim follows. □

Corollary 5.8 (Subexponential Optimization) *Every strongly bounded CLPP can be optimized in randomized subexponential time.*

5.4 Anti-Attractive Switches Needed in Unbounded Systems

Lemma 5.6 shows that in strongly bounded systems anti-attractive switches cannot improve the value. This is no longer true in non-strongly bounded systems. Consider the system (only Max has choices):

$$\begin{aligned} x_1 &\leq -3, \\ x_1 &\leq 2x_2, \\ x_2 &\leq -2, \\ x_2 &\leq x_1, \end{aligned}$$

which is not strongly bounded. Let σ (stable) select the first constraints for $x_{1,2}$, resulting in a finite optimal value. Both switches for x_1 and x_2 look anti-attractive, but making them simultaneously leads to an unbounded improvement of the target function $\max(x_1 + x_2)$. The situation is depicted in Figure 1.

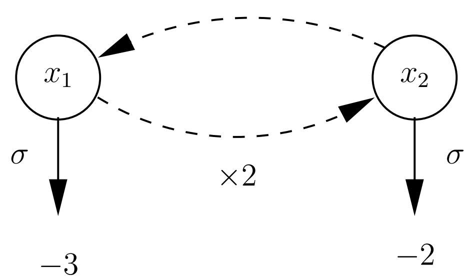

Figure 1: The system above is here shown in graph form. The current strategy σ leads directly to values -3 and -2. Both possible switches are anti-attractive, but making them both would create a loop which makes the system unbounded.

5.5 Weak Duality

Lemma 5.9 (Weak Duality) *For every strongly bounded CLPP instance S and every pair of strategies σ, τ admissible in S, S^d one has:*

$$\max(1^T x | S(\sigma)) \leq \min(1^T x | S^d(\tau)). \tag{3}$$

Proof. We have the following chain of inequalities:

$$\begin{array}{rcll} \max(1^T x|S(\sigma)) & \leq & \max(1^T x|S(\sigma,\tau)) & \text{(for every } \tau,\ S(\sigma,\tau) \text{ is less constrained)} \\ & = & \min(1^T x|S^d(\sigma,\tau)) & \text{(by Proposition 5.3)} \\ & \leq & \min(1^T x|S^d(\tau)) & \text{(for every } \sigma,\ S(\sigma,\tau) \text{ is less constrained).} \end{array}$$

□

5.6 Strong Boundedness is Essential for Weak Duality

The example of the non-strongly bounded system S

$$\begin{array}{rcl} x & \leq & -5 \\ x & \leq & y \\ y & \leq & x \\ y & \leq & 5 \end{array}$$

shows that the strong boundedness requirements in the statement of the Lemma 5.9 is essential. Indeed, when Max chooses $x \leq y$, then Min's best reply is $y \leq 5$, and the optimal solution is $x = y = 5$. But if in the dual system S^d Min chooses $y \geq x$, then Max's best reply will be $x \geq -5$, and the optimal solution is $x = y = -5$. But $5 \not\leq -5$.

This example explains why strong boundedness is essential. In the proof we rely on Proposition 5.3 to conclude that whenever Max and Min stick to σ, τ, the values of the primal and dual systems coincide. In the unbounded case illustrated by the example this is not true: when σ selects $x \leq y$, τ selects $y \leq x$ (both admissible), then $\max(1^T x|S(\sigma,\tau)) = +\infty$, but $\min(1^T x|S^d(\sigma,\tau)) = -\infty$.

5.7 Strong Duality

Theorem 5.10 (Minimax Principle) *For every strongly bounded CLPP instance S one has*

$$\max_\sigma \max(1^T x|S(\sigma)) = \min_\tau \min(1^T x|S^d(\tau)). \tag{4}$$

Proof. Consider a sequence of strategies $\sigma_1, \tau_1, \sigma_2, \tau_2, \ldots, s_p, \tau_p$ corresponding to any sequence of attractive switches for Max in S, where he selects better and better strategies σ_i, and each time Min responds with the best counterstrategy τ_i. The last pair σ_p, τ_p forms a stable pair of strategies. Neither a switch in σ_p can increase, nor in τ_p can decrease the value of $1^T x$. Now, let us start iterative improvement by attractive switches in the dual system S^d starting from the initial strategy τ_p. But this strategy is stable, and by Proposition 5.3 we know that $\max(1^T x|S(\sigma_p,\tau_p)) = \min(1^T x|S^d(\sigma_p,\tau_p))$. Since $\max_\sigma \max(1^T x|S(\sigma)) \geq \max(1^T x|S(\sigma_p,\tau_p)) = \min(1^T x|S^d(\sigma_p,\tau_p) \geq \min_\tau \min(1^T x|S^d(\tau))$, the conclusion follows by Lemma 5.9. □

6 Rational Systems: Boundedness Criteria

For CLPP instances in which all coefficients and constants are rational numbers, we can derive stronger characterizations of boundedness. The rationality assumption is reasonable in most practical circumstances; for instance, it holds for all CLPPs derived from parity, mean payoff, discounted payoff and simple stochastic games (when discounting factors and probabilities are rational).

Lemma 6.1 *For a rational feasible system $Ax \leq b$ the following conditions are equivalent:*

1. $\max(1^T x | Ax \leq b) < +\infty$,
2. $\forall y (Ay \leq 0 \Rightarrow 1^T y \leq 0)$,
3. *The system $Ay \leq 0$, $1^T y > 0$ is infeasible.*

Proof. ($1 \Rightarrow 2$) Assume (1) and suppose, towards a contradiction, that $1^T y' = c > 0$ for some y' satisfying $Ay' \leq 0$. Take any feasible solution x' to $Ax \leq b$ (here we rely on feasibility) and any $\lambda > 0$. Then $x' + \lambda y'$ is feasible for $Ax \leq b$ with the target value larger than $\lambda c > 0$, which can be made arbitrarily large, a contradiction with (1).

($2 \Rightarrow 1$) Assume (2) and suppose, towards a contradiction, that $\max(1^T x | Ax \leq b)$ is unbounded. Then there is an infinite sequence s of integer[3] vectors $Ax_1, Ax_2, \ldots$ such that $Ax_i \leq b$ and $1^T x_i < 1^T x_{i+1}$.

Select an infinite subsequence s_1 of s of vectors with non-increasing first coordinate. Similarly, for $k < n$, select an infinite subsequence s_{k+1} of s_k with non-increasing $k+1$-st coordinate. In the sequence s_n there are two elements $Ax_i \geq Ax_j$ with $i < j$. Consider $x_j - x_i$. By construction, $A(x_j - x_i) \leq 0$, but $1^T(x_j - x_i) > 0$, a contradiction with (2).

(2) is obviously equivalent to (3). □

Note. The feasibility assumption is essential. Consider the infeasible system $\{x \leq y - 1, y \leq x - 1\}$. Then 1 holds, but 2 fails.

Combining Lemma 6.1 with the monotonicity of CLPP constraints, we obtain a stronger characterization of bounded systems.

Corollary 6.2 *Let S be a rational CLPP instance, σ and τ strategies of* MAX *and* MIN, *such that $S(\sigma, \tau)$ written as $Ax \leq b$ is feasible. Then the following conditions are equivalent:*

1. $\max(1^T x | Ax \leq b) < +\infty$,
2. *0 is the unique optimal solution to* $\max(1^T x | Ax \leq 0)$

[3] We may assume that A and b have integer components, by multiplying every constraint by an appropriate integer (least common multiple of denominators). Since $1^T x$ is unbounded on $Ax \leq b$, the feasible region contains integral points x with arbitrarily large values $1^T x$. Since A is integral, it takes integral vector values on such x.

Moreover, both conditions imply that the matrix A is nonsingular.

Proof. ($1 \Rightarrow 2$) Follows from Lemma 4.1.

($2 \Rightarrow 1$.) Follows from Lemma 6.1.

The last claim is implied by (2), since 0 must be the unique solution to $Ax = 0$. □

Remark. Nonsingularity of A does not imply conditions 1, 2 in the lemma above. Consider the system $\{x \leq y, y \leq 2x\}$.

Corollary 6.2 allows us to give an alternative description of what it means to determine if a CLPP is bounded. Given an instance S, is there a strategy σ of Max such that for every strategy τ of Min, the matrix A obtained from $S(\sigma, \tau)$ is such that 0 is the unique optimal solution to $\max(1^T x | Ax \leq 0)$.

References

[1] H. Björklund, O. Nilsson, O. Svensson, and S. Vorobyov. The controlled linear programming problem. Technical Report DIMACS-2004-41, DIMACS: Center for Discrete Mathematics and Theoretical Computer Science, Rutgers University, NJ, September 2004. `http://dimacs.rutgers.edu/TechnicalReports/`.

[2] H. Björklund, S. Sandberg, and S. Vorobyov. A discrete subexponential algorithm for parity games. In H. Alt and M. Habib, editors, *20th International Symposium on Theoretical Aspects of Computer Science, STACS'2003*, volume 2607 of *Lecture Notes in Computer Science*, pages 663–674, Berlin, 2003. Springer-Verlag.

[3] H. Björklund, S. Sandberg, and S. Vorobyov. Randomized subexponential algorithms for parity games. Technical Report 2003-019, Department of Information Technology, Uppsala University, 2003. `http://www.it.uu.se/research/reports/`.

[4] H. Björklund, S. Sandberg, and S. Vorobyov. A combinatorial strongly subexponential strategy improvement algorithm for mean payoff games. In J. Fiala, V. Koubek, and J. Kratochvil, editors, *29th International Symposium on Mathematical Foundations of Computer Science (MFCS'2004)*, volume 3153 of *Lecture Notes in Computer Science*, pages 673–685. Springer, August 2004. Preliminary full version: TR DIMACS-2004-05 `http://dimacs.rutgers.edu/TechnicalReports/`.

[5] A. Condon. On algorithms for simple stochastic games. *DIMACS Series in Discrete Mathematics and Theoretical Computer Science*, 13:51–71, 1993.

[6] V. A. Gurvich, A. V. Karzanov, and L. G. Khachiyan. Cyclic games and an algorithm to find minimax cycle means in directed graphs. *U.S.S.R. Computational Mathematics and Mathematical Physics*, 28(5):85–91, 1988.

[7] M. Jurdziński. Small progress measures for solving parity games. In H. Reichel and S. Tison, editors, *17th STACS*, volume 1770 of *Lect. Notes Comput. Sci.*, pages 290–301. Springer-Verlag, 2000.

[8] W. Ludwig. A subexponential randomized algorithm for the simple stochastic game problem. *Information and Computation*, 117:151–155, 1995.

[9] A. Puri. *Theory of hybrid systems and discrete events systems.* PhD thesis, EECS Univ. Berkeley, 1995.

[10] A. Schrijver. *Theory of Linear and Integer Programming.* John Wiley and Sons, 1986.

[11] J. Vöge and M. Jurdziński. A discrete strategy improvement algorithm for solving parity games. In E. A. Emerson and A. P. Sistla, editors, *CAV'00: Computer-Aided Verification*, volume 1855 of *Lect. Notes Comput. Sci.*, pages 202–215. Springer-Verlag, 2000.

[12] U. Zwick and M. Paterson. The complexity of mean payoff games on graphs. *Theor. Comput. Sci.*, 158:343–359, 1996.