DIMACS Technical Report 2005-20
June 2005

A Subexponential Algorithm for a Subclass of P-Matrix Generalized Linear Complementarity Problems

by

Ola Svensson and Sergei Vorobyov [1,2,3]

[1]DIMACS visitors. Affiliation: Information Technology Department, Uppsala University, 75105 Uppsala, SWEDEN. Correspondence to: `Sergei.Vorobyov@it.uu.se`.
[2]DIMACS hosts: Leonid Khachiyan and Endre Boros.
[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*", "*Interior-Point Methods for Infinite Games*", and "*Infinite Games for Program Verification: Combinatorial Optimization Approach*".

ABSTRACT

We define a nontrivial (polynomially time recognizable, but not known to be polynomial time solvable) subclass of P-matrix Generalized (Vertical) Linear Complementarity Problems, which we call D-matrix GLCPs. D-matrices have nonnegative entries and strict row diagonal dominance. In general, P-matrix LCPs and GLCPs are not known to be polynomial, or even subexponential time solvable, and the property "to be a P-matrix" is coNP-complete.

We describe the first *strongly subexponential* randomized algorithm for D-matrix GLCPs. As a lower bound, we show that simple stochastic games, a long standing problem in $\mathrm{NP} \cap \mathrm{coNP}$, unknown to be polynomial, is subsumed by D-matrix GLCPs.

As an important application, our D-matrix GLCP algorithm gives the best currently available algorithm for computing values of simple stochastic games. At present, the D-matrix GLCP is the only known nontrivial subclass of the P-matrix GLCP solvable in subexponential time.

Contents

1 Introduction

The *Linear Complementarity Problem* (LCP) is a powerful framework for combinatorial and continuous optimization. It has a rich and well-developed theory [18, 11], with numerous important applications. The general problem is NP-hard, but there are many interesting polynomially solvable subclasses, such as Z-matrix and PSD-matrix (positive semidefinite) LCPs [18, 11, 19]. For a prominent class of P-matrix LCPs (possessing unique solutions) there are no currently known polynomial algorithms, but there is strong evidence (NP≠coNP) that the P-matrix LCP is not NP-hard [21]. It is an exciting open problem to invent subexponential algorithms for nontrivial subclasses of P-matrix LCPs.

We consider the *Generalized* LCP (GLCP), first introduced by Cottle and Dantzig [10], also referred to as the *Vertical* LCP in the literature. The GLCP subsumes LCP as a particular case, and it is more flexible and convenient in applications, like game-theoretic ones aimed at in this paper.[1] All definitions of matrix classes (P-, Z-, etc.) extend straightforwardly to the GLCP. To our knowledge, prior to this paper there were no nontrivial, polynomially recognizable and subexponentially solvable classes of P-matrix GLCPs.

In this paper we describe the first nontrivial subclass of P-matrix GLCPs, which we call D-matrix GLCPs or DGLCPs (D- stands for *discounted*), together with a number of randomized subexponential algorithms based on combinatorial linear programming. The class of D-matrices has a simple syntactic description. We show that D-matrix GLCP is nontrivial by subsuming simple stochastic games (currently not known to be polynomial).

Our investigation into the GLCP theory was motivated by applications to solving certain full-information infinite adversary games. In [5] we investigated the LCP approach for expressing and solving *Mean Payoff Games* (MPGs) [14, 15], and developed the first subexponential LCP-based algorithm for MPGs. It was obvious that reductions to LCPs described in [5] applies to wider classes of games, including simple stochastic games and a more general framework of *Controlled Linear Programming Problems* (CLPPs) [2, 1, 3], but the usefulness of such reductions was questionable (prior to this paper) by lack of existing subexponential algorithms for nontrivial classes of (G)LCPs.

Although our LCP-based algorithm for MPGs from [5] is superseded, in principle, by the algorithm presented here (via a known reduction from MPGs to simple stochastic games), the big advantage of the algorithm [5] consists in using a strongly polynomial and simple subroutine for computing shortest paths, whereas the algorithm presented here invokes a more complicated subroutine for solving either a linear program (either nonstrongly polynomial or strongly subexponential), or an instance of the Z-matrix GLCP. As the reduction involves high precision arithmetic with long rationals, this leads to worse hidden constants. In contrast, the direct MPG-LCP algorithm [5] is purely combinatorial. Curiously, we show that MPG-generated LCPs are non-P-matrix, while reduction to simple stochastic games results in D-, hence P-matrix GLCPs. Certainly, the algorithm [5] will become obsolete the day P-matrix GLCPs are proved polynomial.

[1] The advantage of the GLCP is that it allows for more than binary complementarity condition (in contrast to the standard LCP), which is more suitable for describing games on graphs with arbitrary outdegree.

Morris [22] considers randomized pivoting schemes borrowed from combinatorial linear programming for the simple restricted case of square (standard) P-matrix LCPs, corresponding to Boolean hypercubes with edge orientations (we consider Generalized LCPs corresponding to oriented products of simplices). He shows that principal pivoting algorithms may cycle, and in this case the randomized algorithms fair poorly. Morris does not define any syntactic subclasses of P-matrix LCPs solvable in subexponential time (except those for which the induced ordering of the cube is acyclic – a semantic class). In contrast, we define a syntactic, polynomially recognizable, class of vertical block D-matrices (subclass of P-matrices) for which our randomized pivoting is monotonic, hence acyclic, and is guaranteed to be subexponential.

Alternative aproaches to solving simple stochastic and mean payoff games are described in [6, 7].

Outline. The paper is organized as follows. After basic GLCP definitions we introduce the class of D-matrix GLCPs (DGLCPs), show that D-matrices are P-matrices and DGLCPs possess unique solutions. After that we introduce a class of switching algorithms for DGLCPs and prove monotonicity of switching, crucial for termination and subexponential analysis. Then we describe a family of subexponential algorithms and give a reduction from linear minimax systems (subsuming simple stochastic games) to DGLCPs.

2 Generalized Linear Complementarity Problem

We start by reminding the reader of the standard square-matrix LCP [18, 11].

Definition 2.1 (LCP) *The* Linear Complementarity Problem *is specified as follows.*
Given: *a matrix $N \in \mathbb{R}^{k \times k}$ and a vector $q \in \mathbb{R}^k$.*
Find: *vectors $w, z \in \mathbb{R}^k$ satisfying*

$$\begin{aligned} & w = q + Nz, & \\ & w \geq 0, z \geq 0, & \textit{(Nonnegativity)} \\ & z_i w_i = 0, \ \textit{for } i = 1, \ldots, k. & \textit{(Complementarity)} \end{aligned} \tag{1}$$

The complementarity condition can be equivalently expressed as *orthogonality* $z^T w = 0$. Assuming nonnegativity, both say that whenever a component of w is nonzero, the corresponding component of z should be zero, and vice versa.

The Vertical LCP is a generalization of the Linear Complementarity Problem. It was first introduced and studied by Cottle and Dantzig [10] (called Generalized LCP there; we will use both terms as equivalent). The vertical LCP is not limited to square matrices and binary complementarity conditions as in (1). As a consequence, it is more suitable for expressing games played on a graph with arbitrary out-degree, without first reducing them to binary ones; see [3]. Before stating the problem, we will need the definition of a *vertical block matrix.*

Definition 2.2 *A* vertical block matrix of type $(p_1, \ldots, p_k)$ *is a real block matrix*

$$N = \begin{bmatrix} N^1 \\ \vdots \\ N^k \end{bmatrix}$$

with the j-th block N^j of order $p_j \times k$. Thus for $p = \sum_{j=1}^{k} p_j$ the matrix N has order $p \times k$.

Note that in N the number of blocks equals the number of columns.

Here comes the main definition.

Definition 2.3 (GLCP/VLCP) *The* Generalized or Vertical *LCP is specified as follows.*

Given: *a* vertical block matrix N *of type* $(p_1, \ldots, p_k)$ *and a constant vector* q *decomposed in conformity with* N*:*

$$q = \begin{bmatrix} q^1 \\ \vdots \\ q^k \end{bmatrix}$$

Find: *a vector* $w \in \mathbb{R}^p$ *(decomposed as* q*) and* $z \in \mathbb{R}^k$ *satisfying*

$$\begin{aligned} & w = q + Nz, \\ & w \geq 0, z \geq 0, \\ & z_i \textstyle\prod_{j=1}^{p_i} w_j^i = 0, \text{ for } i = 1, \ldots, k, \quad \textit{(Generalized Complementarity)} \end{aligned} \tag{2}$$

where $p = \sum_{i=1}^{k} p_i$. □

The standard LCP (1) is a special case of the Generalized/Vertical LCP (2), with all blocks of size 1 and square matrix N.

Many results of the GLCP depend on the matrix structure. The analysis of the matrix structure of a GLCP, often boils down to the investigation of representative submatrices.

Definition 2.4 (Representative Submatrix) *A square submatrix* M *of a vertical block matrix* N *is called a* representative submatrix *if its* i*-th row is drawn from* N^i*, the* i*-th block of* N. □

The following classes of matrices are well investigated in the literature [18, 11, 19]. Every class is first defined for square matrices and then the definition is extended in a standard way to block matrices by stipulating the property to hold for all representative submatrices.

Definition 2.5 *A square matrix* M *is:*

1. *a P-matrix if all principal minors of* M *are positive;*
2. *a Z-matrix if all off-diagonal elements of* M *are nonpositive; if* M *also is a P-matrix it is called a K-matrix;*
3. *strictly row diagonally dominant if* $|M_{ii}| > \sum_{j \neq i} |M_{ij}|$ *for each row* i*;*
4. *diagonally positive if all diagonal elements of* M *are positive.*

A vertical block matrix is a P-matrix, Z-matrix, etc., if all its representative submatrices are square P-matrices, Z-matrices, etc, respectively. □

The property of being a P-matrix is coNP-complete [12].

3 D-Matrices and Discounted GLCPs

In this paper we introduce a new class of vertical matrices and corresponding GLCPs called *Discounted*, D-matrices and DGLCPs for short. We will demonstrate that D-matrices are P-matrices, and *unique* solutions to DGLCPs can be found in randomized *subexponential* time, which cannot be done (at least not known yet) for general P-matrices [22]. We also show that stopping simple stochastic games (as general as simple stochastic games) are easily expressible as D-matrix GLCPs. Here is our main definition.

Definition 3.1 (Discounted Vertical Matrix, Discounted LCP) *A vertical block matrix N is* Discounted, *or D-matrix, if N is of the form depicted in Figure 1 and has the following properties:*

1. *all elements of N are non-negative;*
2. *every representative submatrix of N has a unit main diagonal;*
3. *the remaining nonzero entries are located in the gray area;*
4. *N is strictly row diagonally dominant.*

A D-matrix GLCP is called Discounted *GLCP, or DGLCP for short.* □

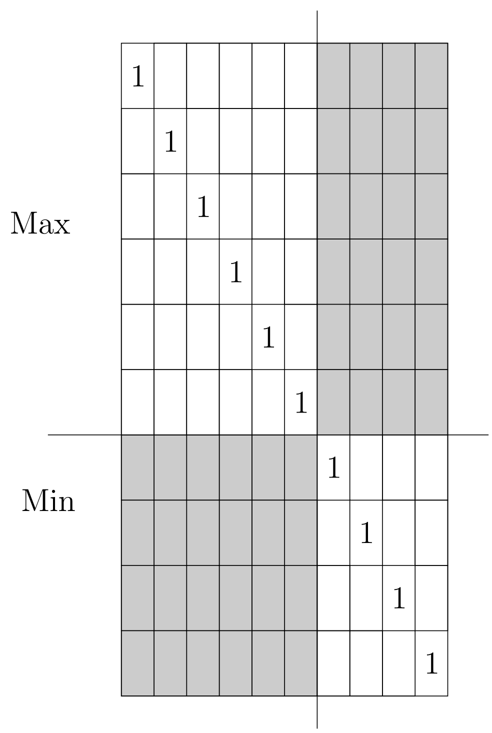

Figure 1: The structure of D-matrix.

Remark. Note that the property of being a D-matrix is easily polynomial time recognizable. Assuming the unit main diagonal is just a matter of technical convenience. Indeed, if other properties are satisfied for N, then multiplying rows of the GLCP $w = q + Nz$ by appropriate positive constants yields an equivalent system $c^T w = w' = q' + N'z$, with a D-matrix N'. We will show that a D-matrix GLCP has a unique solution w', z for each q', which allows for reconstructing the unique solution to $w = q + Nz$. □

Notational Conventions. For the reasons, which will become clear shortly, the upper part of a D-matrix, consisting of $n \leq k$ blocks, is associated with the player Max, and the lower part with the player Min. Thus the range of blocks and columns is split between Max and Min. We will write $i \in \textsc{Max}$ or $j \in \textsc{Min}$ meaning that the corresponding index is in the range of one of the players. By $z_{|\textsc{Min}}$ we will denote vector z with all components of Max replaced with zeros, and similarly for $z_{|\textsc{Max}}$, $w_{|\textsc{Max}}$, $w_{|\textsc{Min}}$.

By U^i_j or L^i_j, respectively, depending on whether $i \in \textsc{Max}$ or $i \in \textsc{Min}$, we will denote the j-th row in the i-th block of a D-matrix, with the 1 in the i-th coordinate (main diagonal) replaced with 0. We will call vectors U^i_j and L^i_j *discounted*, because they have nonnegative coordinates summing up to a number < 1. □

Note that by our conventions $U^i_j\ z_{|\textsc{Min}} = U^i_j\ z$ and $L^i_j\ z_{|\textsc{Max}} = L^i_j\ z$.

Remark. When in Figure 1 one of the partitions Max, Min is *empty*, every representative submatrix is the identity. This case *does not correspond* to one-player games. One player games are captured by Z- and ZD$^+$-matrices; see Section 8. □

4 D-Matrices are P-Matrices

We start by demonstrating that D-matrices form a subclass of P-matrices.

Theorem 4.1 *Every D-matrix is a P-matrix.*

Proof. It is well known that a matrix M is a P-matrix iff M reverses the sign of no nonzero vector [11, Theorem 3.3.4, p. 147] and [18, Theorem 3.12, p. 210], i.e., for every z

$$[z_i(Mz)_i \leq 0 \text{ for all } i] \Rightarrow [z = 0]. \tag{3}$$

Take any *representative submatrix* M of N. Suppose, there exists a nonzero vector z violating (3), i.e., making the premise true and the conclusion false. Identify the component z_i of z with the largest absolute value. Then it is easy to see the contradiction, $z_i(Mz)_i > 0$ (i.e., the premise is false), because M is strictly row diagonally dominant (thus $(Mz)_i$ has the same sign as z_i). Consequently, every *representative submatrix* M of N is a P-matrix and therefore, N is a P-matrix. □

Actually, the same proof gives a stronger (but not needed here)

Corollary 4.2 *Every diagonally positive and strictly row diagonally dominant matrix is a P-matrix.* □

We show in Section 13 that the subclass of D-matrices is *nontrivial*, i.e., there are no currently known polynomial algorithms for D-matrix GLCPs. To our knowledge, this is the first nontrivial subclass of P-matrices for which the GLCP can be solved in subexponential time.

5 Uniqueness of Solutions for Discounted GLCP

It is well known that for every square P-matrix N and column vector q the LCP (1) has a *unique* solution [18, 11]. Cottle&Dantzig [10] show that every P-matrix GLCP (N, q) has a solution[2]. To make the paper self-contained, we present a simple proof that DGLCPs have unique solutions. This result is crucial for our algorithms and greatly simplifies proofs below.

Theorem 5.1 (Uniqueness) *For every D-matrix N and right-hand side vector q the GLCP (N, q) has a unique solution.*

Proof. Suppose, (w, z) and (w', z') are two different solutions to a DGLCP (N, q). Obviously, $z = z'$ implies $w = w'$, because z uniquely determines w by $w = q + Nz$. Therefore, assume (wlog) that $\max_j(z_j^{'} - z_j) \geq \max_j(z_j - z_j^{'})$, and select i for which the difference $z_i' - z_i > 0$ is the largest.

Since $z_i' > 0$, complementarity allows us to select an index j in block i in such a way that

$$\begin{array}{ccccccccc} 0 & \leq & w_j^i & = & q_j^i & + & z_i & + & U_j^i\, z, \\ 0 & = & w_j'^i & = & q_j^i & + & z_i^{'} & + & U_j^i\, z^{'}. \end{array} \tag{4}$$

(the case when the block i corresponds to the constraint $w_j^i = q_j^i + z_i + L_j^i\, z$ in the low part is completely similar.)

Subtracting the first line from the second we get

$$0 \geq z_i^{'} - z_i + U_j^i(z^{'} - z) \text{ or } z_i^{'} - z_i \leq U_j^i(z - z^{'}), \tag{5}$$

which gives an immediate contradiction, because $z, z^{'} \geq 0$, the vector U_j^i is discounted, and by the choice of $z_i^{'} - z_i$ with the largest difference. $\square$

Actually, with almost the same proof, recalling that U_j^i is the the j-th row in block N^i with the i-th element replaced with 0, we obtain the following more general

Theorem 5.2 (Uniqueness) *For every diagonally positive and strictly row diagonally dominant matrix N the GLCP (N, q) has a unique solution for any right-hand side vector q.*

Proof. The modifications in the proof above are as follows. In (4) z_i and $z_i^{'}$ are replaced with αz_i and $\alpha z_i^{'}$ (for some $\alpha > 0$). Consequently, (5) results in $\alpha(z_i^{'} - z_i) \leq U_j^i(z - z^{'})$, or $z_i^{'} - z_i \leq (1/\alpha) \cdot U_j^i(z - z^{'})$, and the vector $(1/\alpha) \cdot U_j^i$ is discounted. $\square$

[2]but at the time of writing we are unaware whether uniqueness holds.

6 Strategies, Attractiveness, Switches, Stability

Informally, a strategy stipulates which of the complementary variables we want to make 0 in a solution.

Definition 6.1 (Strategy) *A strategy σ_i for the block $i \in$ MAX is a selection of either $\sigma_i = \{z_i = 0\}$ or $\sigma_i = \{w^i_j = 0\}$, for some $j \in \{1, \dots, p_i\}$. A full* MAX *strategy σ consists of selections σ_i for all blocks $i \in$ MAX and is denoted $\sigma = \{\sigma_1, \dots, \sigma_n\}$. A partial* MAX *strategy consists of strategy selections for some* MAX *blocks.* □

After selecting a strategy σ_i for block $i \in$ MAX, we have either $z_i = 0$ or $z_i = -q^i_j - U^i_j \cdot z_{|\text{MIN}}$.

Denote by $GLCP_{\sigma_i}(N, q)$ the system (N, q) with the i-th block and i-th column removed and remaining occurrences of z_i replaced with 0 (if $\sigma_i = \{z_i = 0\}$) or with $-q^i_j - U^i_j \cdot z_{|\text{MIN}}$ (if $\sigma_i = \{w^i_j = 0\}$). Note that $GLCP_{\sigma_i}(N, q)$ is not a DGLCP any longer, in general.

Given a MAX strategy σ for all blocks we denote the resulting system, after removing all blocks of MAX and replacing all $z_{|\text{MAX}}$, by $GLCP_\sigma(N, q)$.

Note that after finding a solution $w_{|\text{MIN}}$ and $z_{|\text{MIN}}$ for $GLCP_\sigma(N, q)$, the values $w_{|\text{MAX}}$ and $z_{|\text{MAX}}$ are easily and uniquely calculated by substitution but these values may be *negative*. The fact that some values are negative means that we made a mistake in selecting a strategy and some switches have to be made.

Definition 6.2 (Attractiveness, Switches, and Stability) *For a full* MAX *strategy σ, let w^* and z^* be the complementarity vectors after calculating $w_{|\text{MAX}}$ and $z_{|\text{MAX}}$ from the solution of $GLCP_\sigma(N, q)$, as explained above.*

1. *Say that a pivot to w^i_j or z_i is* attractive*, for $i \in$ MAX, if $w^{*i}_j < 0$ or $z^*_i < 0$, respectively.*

2. *An attractive switch for σ in block $i \in$ MAX results from making an attractive pivot, replacing σ_i with $\sigma'_i = \{w^i_j = 0\}$ or $\sigma'_i = \{z_i = 0\}$.*

3. *The strategy σ is* stable *if w^* and z^* are nonnegative, i.e., give a solution to the DGLCP(N,q).* □

7 Discounted GLCP Resulting After Fixing a Partial Strategy

Making a partial (or complete) substitution of a strategy in a DGLCP results in a GLCP with a unique solution.

Theorem 7.1 *For a $DGLCP(N,q)$ with matrix N of order $p \times k$ and n blocks of* Max*, the resulting $GLCP_{\sigma_{1,\ldots,l}}(N,q)$, for any* Max *strategy $\sigma_{1,\ldots,l}$ ($l \leq n$), has a unique solution.*

Proof. The proof is by induction. The inductive hypothesis (IH) is that the matrix in the system $GLCP_{\sigma_{1,\ldots,l}}(N,q)$ has the following properties.

1. The Max partition is as in the Definition 3.1 of the DGLCP.
2. The Min partition is strictly row diagonally dominant with a positive diagonal ≤ 1 (in every representative matrix).

For the base case $l = 0$, the original system $DGLCP(N,q)$ satisfies the IH by definition.

For the inductive step, assume, IH is true for $l-1$, with $l \leq n$. We want to prove the IH for the $GLCP_{\sigma_{1,\ldots,l}}(N,q)$. Let A be the vertical block matrix for the system $GLCP_{\sigma_{1,\ldots,l-1}}(N,q)$.

Case 1 ($\sigma_l = \{z_l = 0\}$). Removing the l-th column and block from A and substituting z_l with 0 in the remainder results is a matrix $GLCP_{\sigma_{1,\ldots,l}}(N,q)$ obviously satisfying the IH.

Case 2 ($\sigma_l = \{w^l_j = 0\}$). We remove the l-th block and column and substitute z_l with $-U^l_j \cdot z - q^l_j$. The resulting matrix B will be the matrix for the system $GLCP_{\sigma_{1,\ldots,l}}(N,q)$. It is immediate that the Max partition of B satisfies the IH.

Moreover, every Min row in A can be written as

$$w^d_i = A^d_i \cdot z + q^d_i = \lambda_1 z_1 + \ldots \lambda_l z_l + \cdots + \lambda_k z_k + q^d_i$$

By positivity of the diagonal and strict row diagonal dominance (IH), we have

$$\lambda_d > \sum_{x \neq d} |\lambda_x|.$$

After substituting z_l with $-U^l_j \cdot z - q^l_j$ z_l, the same row of B will have the form

$$w^d_i = \lambda_1 z_1 + \cdots - \lambda_l U^l_j \cdot z_{|\text{Min}} + \cdots + \lambda_k z_k + q^d_i - \lambda_l q^l_j.$$

Since $\sum_y |U^l_{jy}| < 1$ implies $|\lambda_l| \sum_y |U^l_{jy}| < |\lambda_l|$, it follows that

$$\lambda_d > \sum_{x \neq d} |\lambda_x| > \sum_{x \neq d, x \neq l} |\lambda_x| + |\lambda_l| \sum_y |U^l_{jy}|.$$

Thus, after the substitution of z_l the row remains strictly diagonally dominant with a positive diagonal. Moreover, the diagonal entry is ≤ 1, because at every step one MAX variable z_i with a positive coefficient is replaced by either 0 or a nonpositive linear polynomial depending on variables z_j, $j \in$ MIN. Therefore, λ_d may only decrease. Consequently, the MIN part of $GLCP_{\sigma_{1,\ldots,l}}(N,q)$ also satisfies the IH.

We have proved that the IH holds for all $GLCP_{\sigma_{1,\ldots,l}}(N,q)$, where $l \leq n$. By Theorem 5.2, the IH implies the uniqueness of a solution. □

8 One-Player Case Yields Discounted Z-Matrices

We have a special restricted subclass of block Z-matrices resulting from substituting full MAX strategies in DGLCPs, which deserve a special name.

Definition 8.1 *A vertical block is called a ZD$^+$-matrix if it is:*

1. *a Z-matrix, i.e., all off-diagonal elements (in the representative matrices) are* ≤ 0*;*
2. *diagonally positive with all diagonal elements (in the representative matrices) in the range* $(0, 1]$*;*
3. *strictly row diagonally dominant.* □

Theorem 8.2 *For every full* MAX *strategy* σ *the matrix in the system* $GLCP_\sigma(N, q)$ *is a ZD$^+$-matrix possessing a unique solution.*

Proof. Recall that the first n blocks belongs to MAX and the remaining m blocks belongs to MIN. Every MIN row in $DGLCP(N, q)$ can be written as

$$w_j^{n+d} = z_{n+d} + \lambda_1 z_1 + \cdots + \lambda_n z_n + q_j^{n+d}, \tag{6}$$

with $\lambda_l \geq 0$. After selecting a strategy σ, either $z_i = 0$ or $z_i = -U_j^i z_{|\mathrm{MIN}} - q_j^i$, for $1 \leq i \leq n$ and some $j \in \{1, \ldots, p_i\}$. Substituting the value for every z_i, $1 \leq i \leq n$, into (6) will result in a nonpositive coefficient in front of z_j, where $n < j \leq k$ and $j \neq n+d$, because U is nonnegative. Hence, all off-diagonal entries of $GLCP_\sigma(N, q)$ will be nonpositive. Consequently, the matrix of $GLCP_\sigma(N, q)$ is a Z-matrix.[3]

The remaining conditions 2, 3 of Definition 8.1 for the matrix of $GLCP_\sigma(N, q)$ and solution uniqueness follow from the inductive proof of Theorem 7.1.[4] □

Remark. Without loss of generality we may assume that ZD$^+$-matrices in the statement of Theorem 8.2 have *ones* on the main diagonals of representative matrices. Indeed, by multiplying equation of $w = q + Nz$ by appropriate positive constants we get an equivalent system $c^T w = q' + N'z$, with coordinate positive c^T and unit diagonal N', which still has a unique solution. After finding this unique solution (after renaming $w' = c^T w$) we uniquely recover the original solution for w. □

Corollary 8.3 *Every ZD$^+$-matrix is a K-matrix.* □

[3] This part of the proof does not use any row diagonally dominance or discountedness properties. It only relies on the bipartite structure, shown in Figure 1, and can therefore be generalized.

[4] This part of the proof depends on diagonal dominance and discountedness.

9 Monotonicity: Attractiveness is Profitable

Monotonicity of attractive switches/pivots (to be explained shortly) is the crucial property ensuring termination of our pivoting algorithms and allowing for subexponential upper bounds.

To simplify the proof (to reduce the number of cases considered), we assume that the algorithms always start from a Max strategy selecting $z_i = 0$ for all $i \in \text{Max}$ and proceed by making attractive switches/pivots.[5]

A subexponential randomized policy is described in Section 11, but monotonicity proved here guarantees finite termination of any sequence of attractive switches/pivots. For our purposes, making just one attractive switch at a time is enough, but one can consider a generalization when several such pivots are made simultaneously.

To simplify notation we make a convention to denote solutions to the GLCPs before and after a switch as non-primed w, z and primed $w^{'}$, $z^{'}$, respectively.

Definition 9.1 *The* value $val(w, z)$ *of a solution* (w, z) *to a DGLCP equals*

$$\sum_{i \in \text{Max}} z_i - \sum_{k \in \text{Min}} z_k. \tag{7}$$

Monotonicity of attractive switches guarantees that this value *strictly monotonically increases*, which immediately follows from the more general

Theorem 9.2 (Monotonicity) *For every attractive switch/pivot in any DGLCP instance from solution* (w, z) *to solution* $(w^{'}, z^{'})$ *the following properties are satisfied:*

1. $z_i^{'} - z_i \geq 0$ *for each* $i \in \text{Max}$ *(monotonic* non-decrease*);*
2. *at least one inequality above is* strict*, namely the one in the block where an attractive switch was made;*
3. $z_k^{'} - z_k \leq 0$ *for each* $k \in \text{Min}$ *(monotonic* non-increase*).*

Proof. Suppose, an attractive switch/pivot in block $i \in \text{Max}$ results in a new strategy with $\sigma_i^{\prime} = \{w_j^{\prime i} = 0\}$.

Let us start by proving Claim 3 by contradiction.[6] Since the switch was attractive, the following constraints are satisfied (the first line means attractiveness, the second stipulates that that after a switch we impose $w_j^{\prime i} = 0$): [7]

$$\begin{array}{ccccccccc} 0 & > & w_j^i & = & q_j^i & + & z_i & + & U_j^i \, z, \\ 0 & = & w_j^{\prime i} & = & q_j^i & + & z_i^{'} & + & U_j^i \, z^{'}. \end{array} \tag{8}$$

[5] With this assumption, every switch away from $z_i = 0$, $i \in \text{Max}$, will be definitive, i.e., the algorithm will never switch back to $z_i = 0$. The extension to an arbitrary initial strategy is pretty straightforward.

[6] This is the only part of the proof that relies on discountedness.

[7] It is not important here whether before the switch σ_i was $\{z_i = 0\}$ or $\{w_{j'}^i = 0\}$.

Suppose, toward a contradiction, that some z_k (for $k \in \textsc{Min}$) increases its value, and select k yielding the *largest* increase $c > 0$,

$$c = z_k^{'} - z_k > 0. \tag{9}$$

Subtracting the first line in (8) from the second one, we get

$$-(z_i^{'} - z_i) < U_j^i(z^{'} - z) \leq \lambda \cdot c,$$

where the last inequality holds for some $0 < \lambda < 1$, because $U_j^i\, z$ is discounted, depends only on variables of $\textsc{Min}$ ($U_j^i z = U_j^i z_{|\textsc{Min}}$), z, $z^{'}$ are nonnegative, and by the choice of k.

Consequently, for the block $i \in \textsc{Max}$ in which the switch was made,

$$z_i^{'} - z_i > -\lambda \cdot c.$$

Similarly, in each block $i \in \textsc{Max}$ in which there was no switch,

$$z_i^{'} - z_i \geq -\lambda \cdot c.$$

(the only difference consists in replacing $>$ in the first line of (8) with $=$, which results in a non-strict inequality.)

Now let us look in the selected block $k \in \textsc{Min}$. For *every* constraint m in this block, before and after the switch, we have, respectively:

$$\begin{array}{rcccccc} w_m^k & = & q_j^k & + & z_k & + & L_j^k\, z \\ w_m^{'k} & = & q_j^k & + & z_k^{'} & + & L_j^k\, z^{'} \end{array} \tag{10}$$

Subtracting the first line of (10) from the second one we get

$$w_m^{'k} - w_m^k = (z_k^{'} - z_k) + L_m^k(z^{'} - z) \geq c - \lambda \cdot c > 0,$$

because $L_m^k\, z$ is a discounted polynomial depending only on variables z_i with $i \in \textsc{Max}$, and for all such we proved $z_i^{'} - z_i \geq -\lambda \cdot c$.

The last chain of inequalities leads to a contradiction. Indeed, $w_m^{'k} - w_m^k > 0$ plus nonnegativity of w_m^k imply $w_m^{'k} > 0$ for *every* m in block k. By assumption $z_k^{'} - z_k > 0$ and nonnegativity of z_k, we also have $z_k^{'} > 0$. But this implies that the complementarity $z_k^{'} \prod_{m=1}^{n_m} w_m^{'k} > 0$ in block k is violated. This shows that the increase (9) for z_k, $k \in \textsc{Min}$, cannot happen, which proves Claim 3.

Let us now prove Claims 1 and 2, which depend on non-negativity of coefficients in U_j^i, but not on discountedness.

Assume the attractive switch happens in block i and consists in switching from $\sigma_i = \{w_l^i = 0\}$ to $\sigma_i' = \{w_j^{'i} = 0\}$, [8] i.e., the following system of constraints is satisfied:

[8]The case not covered here, but completely analogous, is when the switch is made from $\sigma_i = \{z_i = 0\}$, i.e., the first line is replaced with $z_i = 0$.

$$
\begin{array}{ccccccccccc}
0 & = & w_l^i & = & q_l^i & + & z_i & + & U_l^i\, z, \\
0 & > & w_j^i & = & q_j^i & + & z_i & + & U_j^i\, z \\
0 & = & w_j^{'i} & = & q_j^i & + & z_i^{'} & + & U_j^i\, z^{'}
\end{array}
\qquad (11)
$$

(This is consistent with (8); the second line in (11) coincides with the first line in (8). Line 1 expresses the selection before the switch, line 2 attractiveness, and line 3 the selection after the switch.)

From (11) we derive

$$z_i = -q_l^i - U_l^i\, z < -q_j^i - U_j^i\, z \leq -q_j^i - U_j^i\, z^{'} = z_i^{'},$$

where the inequality $\leq$ holds because U_j^i has nonzero (positive) coefficients only for z_k, $k \in \textsc{Min}$ (recall that by our notational convention $U_j^i\, z = U_j^i\, z_{|\textsc{Min}}$), and by the fact that $z_k^{'} - z_k \leq 0$ proved as Claim 3 above.

Therefore, if an attractive switch happened in block i, the corresponding z_i component *strictly increases* $z_i^{'} > z_i$, proving Claim 2.

The block $i \in \textsc{Max}$ in which an attractive switch was *not made* corresponds to the system similar to (11):

$$
\begin{array}{ccccccccccc}
0 & = & w_l^i & = & q_l^i & + & z_i & + & U_l^i\, z, \\
0 & = & w_l^{'i} & = & q_l^i & + & z_i^{'} & + & U_l^i\, z^{'}
\end{array}
\qquad (12)
$$

from which we derive, analogously ($\leq$ holds for the same reason),

$$z_i = -q_l^i - U_l^i\, z \leq -q_l^i - U_l^i\, z^{'} = z_i^{'},$$

which proves Claim 1 $z_i^{'} \geq z_i$ for $i \in \textsc{Max}$ and finishes the proof. □

Remark. It would be interesting to conduct the monotonicity proof without using discountedness, as we did for general CLPPs, relying only on monotonicity (positiveness) of polynomials U_j^i, L_m^k. This is not needed in this paper, but will be useful for generalizations we will publish elsewhere. □

10 Stability Implies Optimality for Discounted GLCPs

The following result is essential for correctness and complexity analysis of our algorithm.

Theorem 10.1 *In every DGLCP instance every stable* Max *strategy determines the same solution.*

(Note that we do not claim stable strategies are unique, they generally are not.)

Proof. Consider any two stable strategies in a DGLCP instance I. Since both are stable, they both determine solutions for I, which should be equal by Uniqueness Theorem 5.1. $\square$

11 Subexponential Algorithms

We now have all the ingredients necessary to describe a class of randomized subexponential algorithms for the D-matrix GLCP, based on combinatorial linear programming schemes due to Kalai [16, 17] and Matoušek-Sharir-Welzl [20].

Given a DGLCP instance (N, q), define a *hyperstructure* as a Cartesian product

$$\mathcal{P} = \prod_{i=1}^{n} S_i, \tag{13}$$

where $n \leq k$ is the number of MAX blocks, k is the total number of blocks, $S_i = \{0, \ldots, p_i\}$, and p_i the size of the i-th block. Intuitively, $\mathcal{P}$ is the space of all MAX strategies, with 0 corresponding to $\sigma_i = \{z_i = 0\}$ and $j > 0$ corresponding to $\sigma_i = \{w^i_j = 0\}$.

Define a *substructure* $\mathcal{P}'$ of $\mathcal{P}$ as a Cartesian product

$$\mathcal{P}' = \prod_{i=1}^{n} S'_i, \tag{14}$$

where $0 \in S'_i \subseteq S_i$ for each $i \in \{1, \ldots, n\}$.

Such a substructure corresponds to the set of strategies in a DGLCP instance (N, q), in which some constraints have been deleted (which remains a DGLCP instance).

Define the *valuation* on the hyperstructure $\mathcal{P}$ as follows. For every MAX strategy $\sigma \in \mathcal{P}$ the $GLCP_\sigma(N, q)$ (obtained by substituting σ considered as an assignment of zeros for z_i and w^i_j, as described in Section 6) is a ZD$^+$-matrix GLCP, possessing a unique solution (w, z) (Theorems 7.1, 8.2). Find this solution as described in Section 12. Assign to σ the value $\nu(\sigma) = val(w, z)$, as defined by (7).

With this valuation,

1. every two neighbors σ and σ' on $\mathcal{P}$ (at Hamming distance 1) corresponding to an attractive switch from σ to σ' have values $\nu(\sigma) < \nu(\sigma')$;

2. on every *substructure* $\mathcal{P}' = \prod_{i=1}^{n} S'_i$ with $0 \in S'_i \subseteq S_i$ for each i, there is a unique stable (optimal) solution/value (cf., Theorem 10.1).

Now numerous well-known randomized subexponential algorithms [16, 17, 20] for finding a (globally) maximal valuation (stable strategy solving the DGLCP instance (N, q)) on the structure $\mathcal{P}$ apply. Roughly (we refer the reader to [16, 17, 20] for details), one of the versions of the algorithm is as follows. Given a structure $\mathcal{P}$, consider the initial strategy $\sigma = \{z_i = 0\}_{i=1,\ldots,n}$, corresponding to the point $\hat{\sigma} = (0, \ldots, 0) \in \mathcal{P}$ (below, for brevity we identify points of hyperstructures with corresponding strategies), and proceed as follows.

1. if $\hat{\sigma} = (\hat{\sigma}_1, \ldots, \hat{\sigma}_n)$ and *the bottom of recursion is hit*, expressed formally as[9]

$$\mathcal{P} = \prod_{i=1}^{n} \{0\} \cup \{\hat{\sigma}_i\}, \tag{15}$$

   then solve an instance of ZD$^+$-matrix $GLCP_\sigma(N, q)$ by one of the algorithms described in Section 12;

2. otherwise, consider a substructure $\mathcal{P}' \subset \mathcal{P}$ containing σ, obtained by (temporarily) throwing away a random $c \in S_i$, $c \neq 0$, $c \neq \sigma_i$ for a random i; [10]

3. apply the algorithm recursively to find a stable (globally maximal) σ^* on $\mathcal{P}'$;

4. return back the last c temporarily thrown away and check whether σ^* is stable in $\mathcal{P}$;

5. if yes, return σ^* as stable (globally maximal) on $\mathcal{P}$;

6. if not, make an attractive switch for σ^*, replacing σ_i^* with c; denote the resulting strategy σ and repeat from step 1.

The analyses of [16, 17, 20] yield the following

Theorem 11.1 *The above algorithm solves an instance of a DGLCP with* n MAX *blocks after expected subexponential* $2^{O(\sqrt{n \log n})}$ *number of switches and invocations of the subroutine solving ZD$^+$-matrix $GLCP_\sigma(N, q)$ in step 1.* □

We show below that ZD$^+$-matrix $GLCP_\sigma(N, q)$ can also be solved in expected subexponential (even *strongly* polynomial) time. This results in the first nontrivial subclass of P-matrix GLCPs solvable in expected subexponential time (in the number of variables); see theorems in Section 12.

Monotonicity of attractive switches (Theorem 9.2) is essential for acyclicity of the algorithm. Uniqueness (Theorems 5.1, 8.2) is crucial for subexponential analysis, because after finding an optimum on a substructure $\mathcal{P}'$ and making the next attractive switch, $\mathcal{P}'$ will never be revisited by the algorithm again (by monotonicity, each attractive switch improves the value), and the subexponential analysis based on *hidden dimensions* applies; see [16, 17, 20] for details.

[9] Technically, we have to keep 0-components in substructures in (15) (cf., (14)) in order to be able to associate elements of hyperstructures to strategies in GLCPs, since, in general, the strategy switch to $\{z_i = 0\}$ is not excluded. However, starting with the strategy $\hat{\sigma} = (0, \ldots, 0)$, and making attractive switches only, by Monotonicity Theorem 9.2, every switch away from $\{z_i = 0\}$ is *definitive*, since z_i for $i \in$ MAX can only increase. Therefore, when (15) holds we immediately know that σ is optimal in $\mathcal{P}$, and it remains to solve the $GLCP_\sigma(N, q)$ to find values to be used in determining further attractive switches.

[10] In other words, we delete a random facet of $\mathcal{P}$ not containing σ.

12 Solving One-Player Z-GLCPs

In the bottom of recursion (when the full MAX strategy σ is fixed) the randomized algorithm described in the previous section solves $GLCP_\sigma(N, q)$, an instance of ZD$^+$-matrix GLCP with a unique solution, as explained in Section 8.

There are several possible algorithms for this problem discussed below.

1. By using the *least element property* [13], solving any feasible Z-matrix GLCP (N, q) amounts to solving a single linear program, minimizing any positive-coordinate linear target function over the feasible domain $\{z : z \geq 0, q + Nz \geq 0\}$. There is a multitude of polynomial (but non-strongly) algorithms for that.

2. The above linear programming problem instance can be solved in randomized strongly subexponential time by the algorithms [16, 17, 20].

   Note that this algorithm is subexponential $2^{O(\sqrt{(k-n)\log(k-n)})}$ in the number $k - n$ of MIN blocks (equals the number of z-variables remaining in $GLCP_\sigma(N, q)$).

3. Reportedly, Venkateswaran [24, 25] invented a *strongly polynomial* algorithm for vertical block Z-matrix GLCP, but at the time of writing we do not have sufficient details about his algorithm.[11] Certainly, it will be much better to use a strongly polynomial subroutine, rather than nonstrongly polynomial or strongly subexponential linear programming algorithms.

The advantage of using options 1 or 2 depends on the size of coefficients in the instance $GLCP_\sigma(N, q)$. Applying option 2, together with Theorem 11.1 results in the following

Theorem 12.1 *A D-matrix GLCP instance with k blocks, n of which belong to* MAX, *can be solved in expected subexponential time* $2^{O(\sqrt{n\log n}+\sqrt{(k-n)\log(k-n)})}$. □

Remarks.

1. With option 3 the bound in Theorem 12.1 reduces to $2^{O(\sqrt{n\log n})}$.

2. The algorithm may be "dualized", making attractive switches in the MIN blocks. This results in the bound $2^{O(\sqrt{m\log m})}$, where $m = \max(n, k-n)$ (if option 2 is used), or $m = \min(n, k-n)$ (if option 3 is used). □

Professor Richard W. Cottle just confirmed [9] that V. Venkateswaran's strongly polynomial algorithm (option 3 above) is well recognized. As result, we have

Theorem 12.2 *A D-matrix GLCP instance with k blocks, n of which belong to* MAX, *can be solved in expected subexponential time* $2^{O(\sqrt{m\log m})}$, *where* $m = \min(n, k-n)$. □

[11] The authors will appreciate any references.

13 Application: MiniMax Systems, Simple Stochastic Games, and Controlled Linear Programming

In this section we describe an application of the theory developed so far to a prominent problem in NP ∩ coNP, unknown to be polynomial . We prove the following

Theorem 13.1 *Every stopping*[12] *simple stochastic game reduces to an instance of the D-matrix GLCP.* □

To our knowledge, the resulting DGLCP-algorithms (Sections 11, 12) are best currently available for simple stochastic games, both asymptotically and regarding the hidden constants in the big-O-notation. Actually, DGLCPs cover a strictly larger than stopping simple stochastic games class of problems, as we explain below.

Let $x = (x_1, \ldots, x_n)$, and $y = (y_1, \ldots, y_m)$ be vectors of variables. Denote by $a^i_j(y)$ and $b^i_j(x)$ linear polynomials[13] with *nonnegative* coefficients and variables in x, y, respectively. Let M be a big enough positive number. [14]

Consider the problem of finding a solution to the linear minimax system

$$\left\{ \begin{array}{lcll} x_i & = & \max & (-M, \quad a^i_1(y), \ldots, a^i_{n_i}(y)) \qquad \text{for } i = 1, \ldots, n, \\ y_i & = & \min & (M, \quad b^i_1(x), \ldots, b^i_{m_i}(x)) \qquad \text{for } i = 1, \ldots, m. \end{array} \right. \tag{16}$$

Obviously, *bipartiteness* (polynomials a^i_j depend on y and b^i_j on x) can be assumed *without loss of generality* here, by introducing auxiliary variables, if necessary. This allows for technical and notational simplifications.

Minimax systems of the form (16) with *discounted*[15] polynomials a^i_j, b^i_l subsume *stopping* simple stochastic games, which approximate with an arbitrary precision *simple stochastic games*; see [8, 26, 2, 1, 3] for definitions and details.

Indeed, without loss of generality, we can assume that a simple stochastic game is "tripartite" i.e., every edge from a MAX vertex leads through an *average* vertex to MIN vertices or sinks (see Figure 2) and vice versa, with MAX and MIN interchanged. Condon [8] describes a way to turn the game into a stopping one by selecting an appropriate $\varepsilon > 0$ (so as the values of the original and resulting games do not differ significantly) and incorporating a new average vertex as shown in Figure 2. [16] The resulting game becomes stopping, and the relation between values of vertices x and $y_1, \ldots, y_l$ can be expressed as $x = a(y_1, \ldots, y_l)$,

[12] Stopping condition may be assumed without loss of generality, as long as approximating game values is concerned; see [8, 26] for definitions and details.

[13] Not necessarily homogeneous, i.e., of the form $c^T x + b$.

[14] Usually explicitly given or computable using the standard technique of obtaining bounds on finite solutions to systems of linear equations [23]. For simple stochastic games M is as small as 1.

[15] I.e., with the sum of nonnegative variable coefficients strictly smaller than 1.

[16] Actually, Condon [8] breaks in this way every edge of the initial game graph, but breaking just all edges leaving MAX and MIN vertices suffices. Also Condon's construction is described for (but not limited to) binary outdegree games.

with a *discounted* polynomial (by virtue of multiplying by $1-\varepsilon$). Collecting over all outgoing edges from x we may write $x = \max(-1, a_1(\bar{y}), \ldots, a_m(\bar{y}))$. The construction for MIN vertices is symmetric and results in $y = \min(1, b_1(\bar{x}), \ldots, b_m(\bar{x}))$.

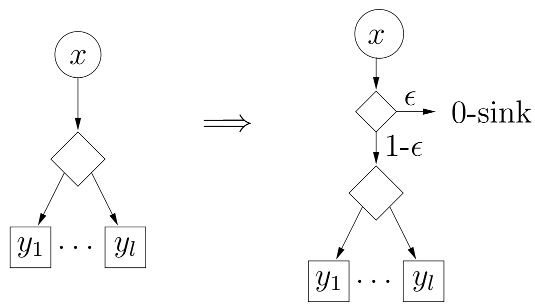

Figure 2: Reduction from a Simple Stochastic Game (SSG) to a stopping one. MAX vertices are round, MIN vertices are square, average vertices are diamond. The sum of probabilities assigned to all edges leaving average vertices sum up to 1.

Now, in a minimax system (16), by introducing nonnegative *slack* variables w_j^i, r_i, $w_j^{'i}$, $r_i^{'}$ we easily get rid of the max and min functions and reduce (16) to solving the following GLCP-like problem (it is not in the canonical form defined in Section 2 yet):

$$\left\{ \begin{array}{lcll} x_i & = & -M + r_i, & \\ x_i & = & a_j^i(y) + w_j^i, & \text{for } i = 1, \ldots, n,\ j = 1, \ldots, n_i, \\ y_i + r_i^{'} & = & M, & \\ y_i + w_j^{'i} & = & b_j^i(x), & \text{for } i = 1, \ldots, m,\ j = 1, \ldots, m_i, \end{array} \right. \tag{17}$$

with the standard complementarity conditions (for all $i = 1, \ldots, n$ and $l = 1, \ldots, m$)

$$w^i,\ r,\ w^{'l},\ r^{'} \geq 0, \quad r_i \prod_{j=1}^{n_i} w_j^i = 0, \quad r_l^{'} \prod_{j=1}^{m_l} w_j^{'l} = 0, \tag{18}$$

which we will not bother to repeat in the sequence of transformations below.

We can now exclude variables x_i and y_i from (17) by substituting $x_i = -M + r_i$ and $y_i = M - r_i^{'}$, which results in the system

$$\left\{ \begin{array}{lcll} -M + r_i & = & a_j^i(\bar{M} - r^{'}) + w_j^i, & \text{for } i = 1, \ldots, n,\ j = 1, \ldots, n_i, \\ M - r_i^{'} + w_j^{'i} & = & b_j^i(-\bar{M} + r), & \text{for } i = 1, \ldots, m,\ j = 1, \ldots, m_i, \end{array} \right. \tag{19}$$

where $\bar{M} - r^{'} = (M - r_1^{'}, \ldots, M - r_m^{'})$ and $-\bar{M} + r = (-M + r_1, \ldots, -M + r_n)$.

Upon rearrangement (19) becomes

$$\left\{ \begin{array}{lcll} w_j^i & = & r_i + a_j^i(r^{'}) - (M + a_j^i(\bar{M})), & \text{for } i = 1, \ldots, n,\ j = 1, \ldots, n_i, \\ w_j^{'i} & = & r_i^{'} + b_j^i(r) - (M + b_j^i(\bar{M})), & \text{for } i = 1, \ldots, m,\ j = 1, \ldots, m_i. \end{array} \right. \tag{20}$$

But this (together with complementarity (18)) is precisely the GLCP with the matrix structure described in Figure 1 of Section 3. It is a D-matrix GLCP, provided that the polynomials a^i_j and b^i_j in (16) were discounted. In this case the unique solution to (20) can be computed by a subexponential algorithm of Section 11. As we mentioned in Section 12, the advantage of the DGLCP approach for solving simple stochastic games consists in using, in the bottom level of recursion, a strongly polynomial algorithm for Z-matrix GLCPs, rather than a nonstrongly polynomial or subexponential linear programming subroutines in alternative approaches [2, 1, 3]. This gives the best currently available algorithm for solving simple stochastic games.

The reduction described also shows that the class of D-matrix GLCPs is *nontrivial*, because polynomial time solvability of simple stochastic games is a long-standing open problem.

The same reduction is applicable to any instance of the recently introduced *Controlled Linear Programming Problem* (CLPP) [2, 1, 3], and gives the best algorithm for *discounted* CLPPs. In cases when a *stable strategy is optimal*, CLPPs are subexponentially solvable by non-LCP means. Some subexponentially solvable subclasses of CLPPs, although reducible to GLCPs, are non yet known to produce subexponentially solvable GLCP instances. This is the case, for example, for *strongly bounded* CLPPs (semantic class). In a forthcoming publication we will present the taxonomy and available CLPP algorithms [4].

In conclusion, consider the case of a one-player stopping simple stochastic game, when MAX (say) has just one choice in every vertex. In this case, rather than applying the general reduction of this section (which produces binary choices $w^i_1 \cdot r_i = 0$), it is better:

1. either to solve a single linear program [8],
2. or, even better, to get a Z-matrix GLCP by substituting every MAX variable x_i with the "only choice" $x_i = a^i_1(y)$; note that this is consistent with letting the slacks $w^i_1 = 0$.

14 Conclusions

We have identified and studied a rich subclass of the P-matrix GLCP called the *D-matrix* or *Discounted* GLCP (DGLCP), with a simple syntactic characterization. We succeeded to define a *monotonic* pivoting/switching rule for DGLCPs, which together with the *solution uniqueness* property of DGLCPs make it possible to apply randomization schemes from combinatorial linear programming. This results in the first *subexponential* algorithm for a nontrivial subclass of P-matrix GLCPs (currently not known to be polynomial time solvable). We demonstrate that the DGLCP subsumes stopping simple stochastic games and a more general *Discounted* Controlled Linear Programming Problem. Furthermore, it turns out that our new DGLCP algorithm is the best known for both problems.

A natural problem for further investigation consists in relaxing the D-matrix restrictions and extending the techniques suggested in this paper to richer subclasses of the P-matrix GLCP, as well as other classes of the GLCP.

Acknowledgement. We are greatful to Professor R. W. Cottle for his communication [9]. Many thanks to Dr. Brenda Latka for suggesting several imprortant corrections and improvements to the presentation.

References

[1] H. Björklund, O. Nilsson, O. Svensson, and S. Vorobyov. Controlled linear programming: Boundedness and duality. Technical Report DIMACS-2004-56, DIMACS: Center for Discrete Mathematics and Theoretical Computer Science, Rutgers University, NJ, December 2004. `http://dimacs.rutgers.edu/TechnicalReports/`.

[2] 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/`.

[3] H. Björklund, O. Svensson, and S. Vorobyov. Controlled linear programming for infinite games. Technical Report DIMACS-2005-13, DIMACS: Center for Discrete Mathematics and Theoretical Computer Science, Rutgers University, NJ, April 2005. `http://dimacs.rutgers.edu/TechnicalReports/`.

[4] H. Björklund, O. Svensson, and S. Vorobyov. Linear complementarity algorithms for controlled linear programming problems. Technical report, DIMACS: Center for Discrete Mathematics and Theoretical Computer Science, Rutgers University, NJ, 2005. to appear.

[5] H. Björklund, O. Svensson, and S. Vorobyov. Linear complementarity algorithms for mean payoff games. Technical Report DIMACS-2005-05, DIMACS: Center for Discrete Mathematics and Theoretical Computer Science, Rutgers University, NJ, February 2005. `http://dimacs.rutgers.edu/TechnicalReports/`.

[6] H. Björklund and S. Vorobyov. A combinatorial strongly subexponential strategy improvement algorithm for mean payoff games. *Discrete Applied Mathematics*, 2005. Accepted for publication, to appear. Preliminary version in MFCS'04, Springer Lecture Notes in Computer Science, vol. 3153, pp. 673-685, and DIMACS TR 2004-05.

[7] H. Björklund and S. Vorobyov. Combinatorial structure and randomized subexponential algorithms for infinite games. *Theoretical Computer Science*, 2005. Accepted for publication, to appear.

[8] A. Condon. The complexity of stochastic games. *Information and Computation*, 96:203–224, 1992.

[9] R. W. Cottle. Private communication. 2005.

[10] R. W. Cottle and G. B. Dantzig. A generalization of the linear complementarity problem. *Journal of Combinatorial Theory*, 8:79–90, 1970.

[11] R. W. Cottle, J.-S. Pang, and R. E. Stone. *The Linear Complementarity Problem.* Academic Press, 1992.

[12] G. Coxson. The P-matrix problem is coNP-complete. *Mathematical Programming*, 64:173–178, 1994.

[13] A. A Ebiefung and M. M. Kostreva. The generalized linear complementarity problem: least element theory and Z-matrices. *Journal of Global Optimization*, 11:151–161, 1997.

[14] A. Ehrenfeucht and J. Mycielski. Positional strategies for mean payoff games. *International Journ. of Game Theory*, 8:109–113, 1979.

[15] 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.

[16] G. Kalai. A subexponential randomized simplex algorithm. In *24th ACM STOC*, pages 475–482, 1992.

[17] G. Kalai. Linear programming, the simplex algorithm and simple polytopes. *Math. Prog. (Ser. B)*, 79:217–234, 1997.

[18] Feng-Tien Yu Katta G. Murty. *Linear Complementarity, Linear and Nonlinear Programming.* Heldermann Verlag, Berlin, 1988. http://ioe.engin.umich.edu/people/fac/books/murty/linear_complementarity_webbook/.

[19] M. Kojima, N. Megiddo, T. Noma, and A. Yoshise. *A Unified Approach to Interior Point Algorithms for Linear Complementarity Problems*, volume 538 of *Lecture Notes in Computer Science*. Springer-Verlag, 1991.

[20] J. Matoušek, M. Sharir, and M. Welzl. A subexponential bound for linear programming. *Algorithmica*, 16:498–516, 1996.

[21] N. Megiddo. A note on the complexity of P-matrix LCP and computing the equilibrium. Technical Report RJ 6439 (62557) 9/19/88, IBM Almaden Research Center, 1988.

[22] W. D. Morris. Randomized pivot algorithms for P-matrix linear complementarity problems. *Mathematical Programming, Ser. A*, 92:285–296, 2002.

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

[24] V. Venkateswaran. An algorithm for the Z-matrix GLCP. In *INFORMS'95*, New Orleans, 1995.

[25] V. Venkateswaran. An algorithm for the block Z-matrix GLCP. In *International Symposium on Mathematical Programming*, page 272, Lausanne, 1997.

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