

Linear Complementarity and P-Matrices for Stochastic Games*

Ola Svensson and Sergei Vorobyov

IDSIA, Lugano, Switzerland and TU Wien, Austria
svorobyov@gmail.com

Abstract. We define the first nontrivial polynomially recognizable subclass of P-matrix Generalized Linear Complementarity Problems (GLCPs) with a subexponential pivot rule. No such classes/rules were previously known. We show that a subclass of Shapley turn-based stochastic games, subsuming Condon's simple stochastic games, is reducible to the new class of GLCPs. Based on this we suggest the new strongly subexponential combinatorial algorithms for these games.

1 Introduction

The *Linear Complementarity Problem* (LCP: find vectors $w, z \geq 0$ satisfying $w = Mz + q$ and $w^T z = 0$ for given real square matrix M and vector q) is a powerful framework for combinatorial and continuous optimization, with a rich theory [21,10] and numerous important applications. The general problem is NP-hard, but there are many rich polynomially solvable subclasses, such as Z-matrix and PSD-matrix (positive semidefinite) LCPs [21,10,17]. For a prominent class of P-matrix LCPs (possessing unique solutions, with positive principal minors of matrix M) there are no currently known polynomial algorithms, but there is strong evidence ($\text{NP} \neq \text{CONP}$) that the P-matrix LCP is not NP-hard [19]. It is an exciting open problem to invent polynomial or at least subexponential algorithms for nontrivial subclasses of P-matrix LCPs [20].

We consider the *Generalized LCP* (GLCP) introduced by Cottle and Dantzig [9], 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.

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

* Supported by the Grant IG2003-2 067 from the Swedish Foundation for International Cooperation in Research and Higher Education (STINT) and the Swedish Research Council (VR) grants “*Infinite Games: Algorithms and Complexity*” and “*Interior-Point Methods for Infinite Games*”.

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

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 Shapley's turn-based stochastic games [22] and Condon's simple stochastic games [8] (both currently not known to be polynomial). To our knowledge, prior to this paper there were no nontrivial, polynomially recognizable and subexponentially solvable classes of P-matrix GLCPs.

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) [13,14], 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 stochastic games and a more general framework of *Controlled Linear Programming Problems* (CLPPs) [3,2,4], but prior to this paper the usefulness of such reductions was questionable by lack of existing subexponential algorithms for nontrivial classes of (G)LCPs. Alternative approaches to solving simple stochastic and mean payoff games are described in [7,6].

Paper Outline. After recalling stochastic games in Section 2, in Section 3 we reduce the value problem for these games to the GLCP problem (basic facts about GLCPs are collected in Appendix A). The structure of the resulting matrices, called D-matrices, is explored in Section 4. D-matrices happen to be P-matrices and D-matrix GLCPs possess unique solutions (Section 5). After that we introduce a class of switching/pivoting algorithms for D-matrix GLCPs (Section 6), analyze the structure of the matrices they produce (Sections 7, 8), prove monotonicity of switching (Section 9) and optimality of stable strategies (Section 10), crucial for termination and subexponential analysis. A family of subexponential algorithms is described in Section 11. Algorithms for solving one-player games at the bottom of recursion are described in Section 12. Missing proofs/details can be found in [23].

2 Shapley's Stochastic Games

For $m \in \mathbb{N}$ denote $[m] = \{i \in \mathbb{N} | 1 \leq i \leq m\}$. In a stochastic game [22] there are finitely many N positions, and players MAX, MIN have finitely many action choices, $[m_k]$, $[n_k]$, respectively, in each position $k \in [N]$. If in position k player MAX selects action $i \in [m_k]$ and MIN simultaneously selects action $j \in [n_k]$, then MAX gets payment a_{ij}^k from MIN, with probability $s_{ij}^k > 0$ the play stops, while with probability $p_{ij}^{kl} \geq 0$ the play proceeds to position l . A particular game Γ^k is obtained by specifying the starting position k . Player MAX wants to maximize, whereas player MIN to minimize the total payoff, which accumulates during the play. Assume, $\sum_{l=1}^N p_{ij}^{kl} = 1 - s_{ij}^k < 1 - s < 1$, $|a_{ij}^k| < M$. Then the probability that a play does not stop after t steps is at most $(1-s)^t$, and the maximal payoff does not exceed $\mathbf{M} = M \sum_{i=0}^{\infty} (1-s)^i = M/s$.

Turn-Based Stochastic Games. Simultaneous move stochastic games thus defined are not perfect information. In *turn-based* stochastic games, for every position k at least one of m_k, n_k equals one. For such a game, let $k \in \text{MAX}$ if $m_k > 1$, $k \in \text{MIN}$ if $n_k > 1$. We call positions k for which both $m_k = n_k = 1$ *unary* and arbitrarily let $k \in \text{MAX}$ or $k \in \text{MIN}$. Turn-based stochastic games are perfect information, solvable in pure positional strategies [22], and the unique value (optimal payoff) for every vertex is determined by the *unique* solution of the system

$$\begin{aligned} v_k &= \max_{i \in [m_k]} (a_{i1}^k + \sum_l p_{i1}^{kl} v_l), \text{ for } k \in \text{MAX}, \\ v_k &= \min_{j \in [n_k]} (a_{1j}^k + \sum_l p_{1j}^{kl} v_l), \text{ for } k \in \text{MIN}. \end{aligned} \tag{1}$$

3 Reducing to Generalized LCP

In this section we show that turn-based Shapley stochastic games are reducible to Generalized LCPs (see Appendix A) of a specific structure.

By introducing, if necessary, auxiliary *unary* positions between positions of the same player, and by appropriately modifying stopping and transitional probabilities, we may assume, with no loss of generality, that the game is *bipartite*, i.e., $p_{ij}^{kl} > 0$ implies $k \in \text{MAX}$ and $l \in \text{MIN}$ or $k \in \text{MIN}$ and $l \in \text{MAX}$.² System (1) can be equivalently presented as:

$$\begin{aligned} v_k &= \max\{ -\mathbf{M}, a_{i1}^k + \sum_l p_{i1}^{kl} u_l \mid i \in [m_k] \}, \text{ for } k \in \text{MAX}, \\ u_k &= \min\{ \mathbf{M}, a_{1j}^k + \sum_l p_{1j}^{kl} v_l \mid j \in [n_k] \}, \text{ for } k \in \text{MIN}, \end{aligned} \tag{2}$$

where we reflect bipartiteness by using v_i/u_i for MAX/MIN variables.

Let us introduce $m_k + 1$ fresh auxiliary *nonnegative* variables $z_k, w_1^k, \dots, w_{m_k}^k \geq 0$ for each variable $v_k \in \text{MAX}$, $n_k + 1$ auxiliary *nonnegative* variables $z_k, w_1^k, \dots, w_{n_k}^k \geq 0$ for each variable $v_k \in \text{MIN}$, and rewrite the system (2) as

$$\begin{aligned} v_k &= z_k - \mathbf{M}, && \text{for } k \in \text{MAX}, \\ v_k &= w_i^k + a_{i1}^k + \sum_l p_{i1}^{kl} u_l, && \text{for } k \in \text{MAX}, i \in [m_k], \\ u_k + z_k &= \mathbf{M}, && \text{for } k \in \text{MIN}, \\ u_k + w_i^k &= a_{i1}^k + \sum_l p_{i1}^{kl} v_l, && \text{for } k \in \text{MIN}, i \in [n_k], \end{aligned} \tag{3}$$

additionally stipulating *complementarity*, i.e.,

$$z_k \cdot \prod_i w_i^k = 0 \text{ for each } k \in [N]. \tag{4}$$

² Although this may blow up quadratically the number of positions, unary positions introduced do not make worse the resulting complexity.

Excluding variables v_i, u_i from (3), we can rewrite it as

$$\begin{aligned} w_i^k &= z_k + P_i^k(\bar{z}_{|\text{MIN}}), \text{ for } k \in \text{MAX}, i \in [m_k], \\ w_i^k &= z_k + P_i^k(\bar{z}_{|\text{MAX}}), \text{ for } k \in \text{MIN}, i \in [n_k], \end{aligned} \tag{5}$$

where: 1) polynomials $P_i^k(\bar{z}_{|\text{MIN}})$ contain only variables z_j for $j \in \text{MIN}$, 2) polynomials $P_i^k(\bar{z}_{|\text{MAX}})$ contain only variables z_j for $j \in \text{MAX}$, 3) these polynomials have all variable coefficients nonnegative, summing up to < 1 (call such polynomials *discounted*). Note that in obtaining this form of system (5) we essentially use bipartiteness, which guarantees that variables w_i^k and z_k appear with non-negative coefficients on *different* sides of equations.

4 D-Matrices and Discounted GLCPs

Finding nonnegative values z_k, w_i^k satisfying (5) and (4) is a well known Generalized (or Vertical) LCP [9]; see Appendix A for a reminder of the main definitions. In this paper, motivated by the special structure of the system (5), 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 [20]. Here is our main

Definition 1 (Discounted Vertical Matrix, Discounted LCP). *A vertical block matrix A is Discounted, or D-matrix, if A is of the form depicted in Figure 1 and has the following properties: 1) all elements of A are non-negative; 2) every representative submatrix of A has a unit main diagonal; 3) the remaining nonzero entries are located in the gray area; 4) A is strictly row diagonally dominant. A D-matrix GLCP is called Discounted GLCP, or DGLCP for short. \square*




Fig. 1. The structure of a D-matrix. 1's denote column vectors of ones.

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 and may be dropped.

The reduction from the previous section shows that the class of DGLCPs is nontrivial:

Theorem 1. *Turn-based Shapley stochastic games and simple stochastic games reduce to the DGLCP.* □

Notational Conventions. From now on we assume that the vertical block matrices considered are of dimension $m \times k$ with $m \geq k$. 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 $i \in \text{MAX}$ or $j \in \text{MIN}$ meaning that the corresponding index is in the range of one of the players. By $z_{|\text{MIN}}$ we will denote vector z with all components of MAX replaced with zeros, and similarly for $z_{|\text{MAX}}$, $w_{|\text{MAX}}$, $w_{|\text{MIN}}$.

By U_j^i or L_j^i , respectively, depending on whether $i \in \text{MAX}$ or $i \in \text{MIN}$, 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_j^i and L_j^i *discounted*, because they have nonnegative coordinates summing up to a number < 1 . Note that by our conventions $U_j^i z_{|\text{MIN}} = U_j^i z$ and $L_j^i z_{|\text{MAX}} = L_j^i z$.

5 D-Matrices Are P-Matrices, Uniqueness of Solutions

We start by an important property that D-matrices form a subclass of P-matrices.

Theorem 2. *Every D-matrix is a P-matrix.*

Proof. Since every representative submatrix of a D-matrix has positive diagonal and is strictly diagonally dominant, it is a P-matrix [10, p. 152]. □

P-matrix GLCPs have unique solutions [24], therefore,

Theorem 3 (Uniqueness). *Every D-matrix GLCP has a unique solution.* □

6 Strategies, Attractiveness, Switches, Stability

A strategy prescribes which of the complementary variables are to be zeroed.

Definition 2 (Strategy). *A strategy σ_i for the block $i \in \text{MAX}$ is a selection of either $\sigma_i = \{z_i = 0\}$ or $\sigma_i = \{w_j^i = 0\}$, for some $j \in \{1, \dots, p_i\}$. A full MAX strategy σ consists of selections σ_i for all blocks $i \in \text{MAX}$ 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 $i \in \text{MAX}$, we have either $z_i = 0$ or $z_i = -q_j^i - U_j^i \cdot z_{|\text{MIN}}$. Denote by $GLCP_{\sigma_i}(A, q)$ the system (A, 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 $-q_j^i - U_j^i \cdot z_{|\text{MIN}}$ (if $\sigma_i = \{w_j^i = 0\}$). Note that $GLCP_{\sigma_i}(A, q)$ 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 $GLCP_{\sigma}(A, q)$. Note that after finding a solution $w_{|\text{MIN}}$ and $z_{|\text{MIN}}$ for $GLCP_{\sigma}(A, q)$, 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 3 (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 $GLCP_{\sigma}(A, q)$, as explained above.*

1. *Say that a pivot to w_j^i or z_i is attractive, for $i \in \text{MAX}$, if $w_j^{*i} < 0$ or $z_i^* < 0$.*
2. *An attractive switch for σ in block $i \in \text{MAX}$ results from making an attractive pivot, replacing σ_i with $\sigma'_i = \{w_j^i = 0\}$ 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(A, q).*

7 GLCPs Resulting from Fixing Partial Strategies

Making a partial (or complete) substitution of a strategy in a DGLCP results in a GLCP with a unique solution, which is a critical invariant for our algorithms:

Theorem 4. *For a DGLCP(A, q) with matrix A of order $p \times k$ and n blocks of MAX, the resulting $GLCP_{\sigma_1, \dots, l}(A, q)$, for any MAX strategy σ_1, \dots, l ($l \leq n$), has a unique solution.*

Proof. By induction. The inductive hypothesis (IH) is that the matrix in the system $GLCP_{\sigma_1, \dots, l}(A, q)$ has the following properties. 1) The MAX partition is as in the Definition 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 $DGLCP(A, q)$ 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 $GLCP_{\sigma_1, \dots, l}(A, q)$. Let A be the vertical block matrix for the system $GLCP_{\sigma_1, \dots, l-1}(A, q)$.

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 $GLCP_{\sigma_1, \dots, l}(A, q)$ obviously satisfying the IH.

Case 2 ($\sigma_l = \{w_j^l = 0\}$). We remove the l -th block and column and substitute z_l with $-U_j^l \cdot z - q_j^l$. The resulting matrix B will be the matrix for the system

$GLCP_{\sigma_1, \dots, l}(A, q)$. It is immediate that the MAX partition of B satisfies the IH. Moreover, every MIN row in A can be written as $w_i^d = A_i^d \cdot z + q_i^d = \lambda_1 z_1 + \dots + \lambda_l z_l + \dots + \lambda_k z_k + q_i^d$.

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 $-U_j^l \cdot z - q_j^l z_l$, the same row of B will be

$$w_i^d = \lambda_1 z_1 + \dots - \lambda_l U_j^l \cdot z_{|\text{MIN}} + \dots + \lambda_k z_k + q_i^d - \lambda_l q_j^l.$$

Since $\sum_y |U_{jy}^l| < 1$ implies $|\lambda_l| \sum_y |U_{jy}^l| < |\lambda_l|$, it follows that $\lambda_d > \sum_{x \neq d} |\lambda_x| > \sum_{x \neq d, x \neq l} |\lambda_x| + |\lambda_l| \sum_y |U_{jy}^l|$. Thus, after the substitution of z_l the row remains strictly diagonally dominant with a positive diagonal. Also, 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 $z_j, j \in \text{MIN}$. Therefore, λ_d may only decrease. Consequently, the MIN part of $GLCP_{\sigma_1, \dots, l}(A, q)$ also satisfies the IH. We have proved that the IH holds for all $GLCP_{\sigma_1, \dots, l}(A, q)$ for $l \leq n$. The uniqueness of their solutions follows from the IH by the proof of Theorem 2 and by Theorem 3. \square

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 4. 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 range $(0, 1]$;
3. strictly row diagonally dominant. \square

Theorem 5. For every full MAX strategy σ the matrix in the system $GLCP_{\sigma}(A, q)$ 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 $DGLCP(A, q)$ can be written as

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

with $\lambda_l \geq 0$. After selecting a strategy σ , either $z_i = 0$ or $z_i = -U_j^i z_{|\text{MIN}} - q_j^i$, for $1 \leq i \leq n$ and some $j \in \{1, \dots, p_i\}$. 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 $GLCP_{\sigma}(A, q)$ will be nonpositive. Consequently, the matrix of $GLCP_{\sigma}(A, q)$ is a Z-matrix.³ The remaining conditions 2, 3 of Definition 4 for

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.

the matrix of $GLCP_\sigma(A, q)$ and solution uniqueness follow from the inductive proof of Theorem 4.⁴ □

Corollary 1. *Every ZD^+ -matrix is a K -matrix.* □

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' .

Definition 5. *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 6 (Monotonicity). *For every attractive switch/pivot in any DGLCP instance from solution (w, z) to solution (w', z') one has:*

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 $\sigma'_i = \{w'_j = 0\}$.

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 $w'_j = 0$):⁷

$$\begin{aligned} 0 &> w_j^i = q_j^i + z_i + U_j^i z, \\ 0 &= w'_j{}^i = q_j^i + z'_i + U_j^i z'. \end{aligned} \tag{8}$$

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

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 $k \in \text{MIN}$) 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 MIN ($U_j^i z = U_j^i z_{|\text{MIN}}$), z, z' are nonnegative, and by the choice of k . Consequently, for the block $i \in \text{MAX}$ in which the switch was made, $z'_i - z_i > -\lambda \cdot c$. Similarly, in each block $i \in \text{MAX}$ 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 $k \in \text{MIN}$. For *every* constraint m in this block, before and after the switch, we have:

$$\begin{aligned} w_m^k &= q_j^k + z_k + L_j^k z \\ w_m'^k &= q_j^k + z'_k + L_j^k z' \end{aligned} \tag{10}$$

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 $i \in \text{MAX}$, 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 $z_k, k \in \text{MIN}$, 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.:

$$\begin{aligned} 0 &= w_l^i = q_l^i + z_i + U_l^i z, \\ 0 &> w_j^i = q_j^i + z_i + U_j^i z, \\ 0 &= w_j'^i = q_j^i + z'_i + U_j^i z'. \end{aligned} \tag{11}$$

(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 $z_k, k \in \text{MIN}$ (recall that by our notational convention $U_j^i z = U_j^i z_{|\text{MIN}}$), 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.

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 $z_i = 0$.

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

$$\begin{aligned} 0 &= w_l^i = q_l^i + z_i + U_l^i z, \\ 0 &= w_l^i = q_l^i + z'_i + U_l^i z' \end{aligned} \tag{12}$$

from which we derive, analogously (\leq holds for the same reason) that $z_i = -q_l^i - U_l^i z \leq -q_l^i - U_l^i z' = z'_i$, which proves Claim 1 $z'_i \geq z_i$ for $i \in \text{MAX}$ and finishes the proof. \square

10 Stability Implies Optimality for Discounted GLCPs

The following result is essential for correctness and complexity analysis of our algorithm. (We do not claim stable strategies are unique, they are generally not.)

Theorem 7. *In every DGLCP instance every stable MAX strategy determines the same solution.*

Proof. Consider any two stable strategies in a DGLCP instance I . Since both are stable, they both determine solutions for I , which are equal by Theorem 3. \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 [15,16] and Matoušek-Sharir-Welzl [18].

Given a DGLCP instance (A, q) , define a *hyperstructure* as a Cartesian product $\mathcal{P} = \prod_{i=1}^n S_i$, where $n \leq k$ is the number of MAX blocks, k is the total number of blocks, $S_i = \{0, \dots, p_i\}$, 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 $\sigma_i = \{w_j^i = 0\}$. Define a *substructure* \mathcal{P}' of \mathcal{P} as a Cartesian product $\mathcal{P}' = \prod_{i=1}^n S'_i$, where $0 \in S'_i \subseteq S_i$ for each $i \in \{1, \dots, n\}$. It corresponds to the set of strategies in a DGLCP instance (A, 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 $GLCP_\sigma(A, q)$ (obtained by substituting σ considered as an assignment of zeros for z_i and w_j^i , as described in Section 6) is a ZD^+ -matrix GLCP, possessing a unique solution (w, z) (Theorems 4, 5). Find this solution as described in Section 12. Assign to σ the value $\nu(\sigma) = \text{val}(w, z)$, 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 7).

Now numerous well-known randomized subexponential algorithms [15,16,18] for finding a (globally) maximal valuation (stable strategy solving the DGLCP instance (A, q)) on the structure \mathcal{P} apply. Roughly (we refer the reader to [15,16,18] for details), one of the versions of the algorithm is as follows. Given a structure \mathcal{P} , consider the initial strategy $\sigma = \{z_i = 0\}_{i=1, \dots, n}$, corresponding to the point $\hat{\sigma} = (0, \dots, 0) \in \mathcal{P}$ (below, for brevity we identify points of hyperstructures with corresponding strategies)⁹, and proceed as follows.

1. if $\hat{\sigma} = (\hat{\sigma}_1, \dots, \hat{\sigma}_n)$ and *the bottom of recursion is hit*, expressed formally as¹⁰

$$\mathcal{P} = \prod_{i=1}^n \{0\} \cup \{\hat{\sigma}_i\}, \tag{13}$$

then solve an instance of ZD⁺-matrix $GLCP_{\sigma}(A, q)$; see 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 ; ¹¹
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 [15,16,18] yield the following

Theorem 8. *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 $GLCP_{\sigma}(A, q)$ in step 1. \square*

In Section 12 we show that ZD⁺-matrix $GLCP_{\sigma}(A, q)$ can also be solved in expected subexponential (or in weakly polynomial) time. This results in the first nontrivial subclass of P-matrix GLCPs solvable in expected subexponential time (in the number of variables).

⁹ For efficiency reasons, it is better to select, in the initial strategy, components $\sigma_l = \{w_l^l = 0\}$ for each unary position l . This neutralizes the effect of introducing many unary positions when reducing to the bipartite case. The correctness of this initial setting is explained by the fact that no switches will ever be made/become attractive in any run of the algorithm.

¹⁰ Technically, we have to keep 0-components in substructures in (13) 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 $\hat{\sigma} = (0, \dots, 0)$, and making attractive switches only, by Monotonicity Theorem 6, every switch away from $\{z_i = 0\}$ is *definitive*, since z_i for $i \in \text{MAX}$ can only increase. Therefore, when (13) holds we immediately know that σ is optimal in \mathcal{P} , and it remains to solve the $GLCP_{\sigma}(A, q)$ to find values to be used in determining further attractive switches.

¹¹ In other words, we delete a random facet of \mathcal{P} not containing σ .

Monotonicity of attractive switches (Theorem 6) is essential for acyclicity of the algorithm. Uniqueness (Theorems 3, 5) 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 [15,16,18] for details.

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 $GLCP_\sigma(A, q)$, an instance of ZD^+ -matrix GLCP with a unique solution, as explained in Section 8. There are several possible algorithms for this problem.

1) By using the *least element property* [12], solving any feasible Z-matrix GLCP (A, q) amounts to solving a single linear program, minimizing any positive-coordinate linear target function over the feasible domain $\{z : z \geq 0, q + Az \geq 0\}$. 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 [15,16,18]. 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 $GLCP_\sigma(A, q)$). The advantage of using options 1 or 2 depends on the size of coefficients in the instance $GLCP_\sigma(A, q)$. Applying option 2, together with Theorem 8 results in

Theorem 9. *A D-matrix GLCP instance with k blocks, n of which belong to MAX, can be solved in expected subexponential time $2^{O(\sqrt{n}\log n + \sqrt{(k-n)\log(k-n)})}$.*

Further improvement will be achieved if a more efficient, strongly polynomial, algorithm for solving ZD^+ -matrix GLCP instances is used at the bottom of recursion. This subject is outside the scope of this paper, deserves a separate careful treatment, and the progress will be reported elsewhere; see, e.g., [1].

13 Conclusions

We identified the first nontrivial subclass of P-matrix Generalized LCPs, which is: 1) polynomial time recognizable (in general, the P-matrix property is coNP-complete); 2) has a very simple syntactical structure; 3) subsumes Shapley's turn-based stochastic games and Condon's simple stochastic games (currently not known to be polynomial time solvable). We suggested the first subexponential pivot rule and algorithm for this class of GLCPs; no such rules were previously known, all were either polynomial or exponential. The resulting algorithm for stochastic games has the same asymptotic behavior as other best currently available algorithms for the problem [4,6].

References

1. D. Andersson and S. Vorobyov. Fast algorithms for monotonic discounted linear programs with two variables per inequality. Manuscript submitted to *Theoretical Computer Science*, July 2006. Preliminary version available as Isaac Newton Institute Preprint NI06019-LAA.
2. H. Björklund, O. Nilsson, O. Svensson, and S. Vorobyov. Controlled linear programming: Boundedness and duality. TR DIMACS-2004-56, Center for Discrete Mathematics and Theoretical Computer Science, Rutgers University, NJ, 2004.
3. H. Björklund, O. Nilsson, O. Svensson, and S. Vorobyov. The controlled linear programming problem. TR DIMACS-2004-41, Center for Discrete Mathematics and Theoretical Computer Science, Rutgers University, NJ, 2004.
4. H. Björklund, O. Svensson, and S. Vorobyov. Controlled linear programming for infinite games. TR DIMACS-2005-13, Center for Discrete Mathematics and Theoretical Computer Science, Rutgers University, NJ, 2005.
5. H. Björklund, O. Svensson, and S. Vorobyov. Linear complementarity algorithms for mean payoff games. TR DIMACS-2005-05, Center for Discrete Mathematics and Theoretical Computer Science, Rutgers University, NJ, 2005.
6. H. Björklund and S. Vorobyov. Combinatorial structure and randomized subexponential algorithms for infinite games. *Theoretical Computer Science*, 349(3): 347–360, 2005.
7. H. Björklund and S. Vorobyov. A combinatorial strongly subexponential strategy improvement algorithm for mean payoff games. *Discrete Applied Mathematics*, 2006, to appear. Available from <http://www.sciencedirect.com/>, 27 June, 2006.
8. A. Condon. The complexity of stochastic games. *Information and Computation*, 96:203–224, 1992.
9. R. W. Cottle and G. B. Dantzig. A generalization of the linear complementarity problem. *Journal of Combinatorial Theory*, 8:79–90, 1970.
10. R. W. Cottle, J.-S. Pang, and R. E. Stone. *The Linear Complementarity Problem*. Academic Press, 1992.
11. G. Coxson. The P-matrix problem is coNP-complete. *Mathematical Programming*, 64:173–178, 1994.
12. 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.
13. A. Ehrenfeucht and J. Mycielski. Positional strategies for mean payoff games. *International Journ. of Game Theory*, 8:109–113, 1979.
14. 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.
15. G. Kalai. A subexponential randomized simplex algorithm. In *24th ACM STOC*, pages 475–482, 1992.
16. G. Kalai. Linear programming, the simplex algorithm and simple polytopes. *Math. Prog. (Ser. B)*, 79:217–234, 1997.
17. 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.
18. J. Matoušek, M. Sharir, and M. Welzl. A subexponential bound for linear programming. *Algorithmica*, 16:498–516, 1996.

19. 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.
20. W. D. Morris. Randomized pivot algorithms for P-matrix linear complementarity problems. *Mathematical Programming, Ser. A*, 92:285–296, 2002.
21. K. G. Murty and F.-T. Yu. *Linear Complementarity, Linear and Nonlinear Programming*. Heldermann Verlag, Berlin, 1988.
http://ioe.engin.umich.edu/people/fac/books/murty/linear_complementarity_webbook/.
22. L. S. Shapley. Stochastic games. *Proc. Natl. Acad. Sci. U.S.A.*, 39:1095–1100, 1953.
23. O. Svensson and S. Vorobyov. A subexponential algorithm for a subclass of P-matrix generalized linear complementarity problems. TR DIMACS-2005-20, Center for Discrete Mathematics and Theoretical Computer Science, Rutgers University, NJ, 2005.
24. B. P. Szanc. *The Generalized Complementarity Problem*. PhD thesis, Rensselaer Polytechnic Institute, Troy, New York, 1989.

A Appendix: Generalized LCP

Definition 6. A vertical block matrix of type (p_1, \dots, p_k) is a real block matrix A of order $p \times k$, where $p = \sum_{j=1}^k p_j$, partitioned by horizontal cuts in blocks A_j of order $p_j \times k$, for $j = 1, \dots, k$. □

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

Here comes the main definition.

Definition 7 (GLCP/VLCP). An instance of the Generalized or Vertical LCP is specified as follows.

Given: a vertical block matrix A of type (p_1, \dots, p_k) and a constant vector q decomposed in conformity with A :

$$q = \begin{bmatrix} q^1 \\ \vdots \\ q^k \end{bmatrix}, \quad A = \begin{bmatrix} A^1 \\ \vdots \\ A^k \end{bmatrix}.$$

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

$$\begin{aligned} w &= q + Az, \\ w &\geq 0, z \geq 0, \\ z_i \prod_{j=1}^{p_i} w_j^i &= 0, \text{ for } i = 1, \dots, k, \text{ (Generalized Complementarity)} \end{aligned} \tag{14}$$

where $p = \sum_{i=1}^k p_i$. □

The standard LCP is a special case of the GLCP (14), with all blocks of size 1 and square matrix A . 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 8 (Representative Submatrix). *A square submatrix M of a vertical block matrix A is called a representative submatrix if its i -th row is drawn from A^i , the i -th block of A . \square*

The following classes of matrices are well investigated in the literature [21,10,17]. 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 9. *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. \square

The property of being a P-matrix is coNP-complete [11].