Efficient Computation of a Canonical Form for a Matrix with the generalized P-property

Walter D. Morris, Jr.
George Mason University
wmorris@gmu.edu

Abstract

We use recent results on algorithms for Markov decision problems to show that a canonical form for a matrix with the generalized P-property can be computed, in some important cases, by a strongly polynomial algorithm.

Key words. Markov Decision Problem, Polytope, Linear Programming, P-matrix

1 Introduction

Suppose that the $m \times n$ real matrix A has the block form

$$A = \left[\ A_1 \ \ | \ \ A_2 \ \ | \cdots \ \ | \ \ A_m\ \right],$$

with each column of block A_j being of the form $e_j - \gamma p_{jk}$, with e_j the j^{th} standard basis vector, γ a scalar in the interval $[0, 1)$, and p_{jk} a vector of nonnegative entries that sum to one. A *Markov decision problem*, as described in [15], is a linear program of the form: maximize $v^T b$, subject to $v^T A \le c^T$ where $b \in \mathbb{R}^m$ is positive and $c \in \mathbb{R}^n$. Two recent papers, [15] and subsequently [7], advanced the theory of such problems significantly. Ye proved that the simplex method, with Dantzig's rule for the entering variable, is a strongly polynomial algorithm to solve the Markov decision problem if the discount factor γ is fixed. Such a positive result for complexity of a pivoting algorithm is rare. The paper of Hansen, Miltersen and Zwick showed a similar result for a considerable generalization of the Markov decision problem, in which one seeks a solution to the optimality conditions of the problem, with some inequalities reversed.

It is natural to seek applications of the advances achieved in these recent papers. The sequence of papers by Kitihara and Mizuno, of which [8] is one example, generalizes Ye's results to other linear programs.

The matrix A from a Markov decision problem has the *P-property*, which means that all square submatrices of A formed by taking one column from each block have determinants of the same non-zero sign. We consider the problem of satisfying the optimality conditions of the Markov decision problem for more general A with the P-property. This problem is known as the generalized linear complementarity problem (LCP) (see [2], [6]). We introduce this problem in Section 2. Theorem 1 of that section describes a canonical form for A with the P-property, introduced in [10] to describe the set of "n-step vectors" of an LCP. One of our goals is to emphasize the significance of this form. Theorem 2 describes a subclass of matrices with the P-property for which the generalized LCP can be solved by solving a linear program. This description is in terms of the canonical form of Theorem 1. The canonical form provides a geometric description of the possible right hand sides of such linear programs if they exist. Theorem 2 also provides a characterization of hidden-K matrices in terms of the canonical form. Theorems 1 and 2 are not essentially new, but the proofs supplied are more streamlined than earlier proofs.

Section 3 discusses the Markov decision problem. The complexity results of [15] and [7] have the term $1 - \gamma$ in the denominator. This leads us to look for equivalent Markov decision problems for which this term is as large as possible. A linear program to find such an optimal problem is formulated. We call this linear program $LP(A)$.

Section 4 presents a two step algorithm to solve $LP(A)$. This method finds the canonical form of Theorem 1 along the way. Section 5 shows that the two step algorithm to solve $LP(A)$ is a strongly polynomial time algorithm, if the optimal value of $LP(A)$ is a fixed positive number.

2 The P-property

Suppose that the matrix $A \in \mathbb{R}^{m\times n}$ has its column set indexed by the set of pairs

$$\{(j,k) : j = 1,2,\ldots,m, k = 1,2,\ldots,n_j\},$$

where $n_1 + n_2 + \cdots + n_m = n$. A *representative* submatrix of A is an $m \times m$ submatrix for which, for each $j = 1,2,\ldots,m$ the j^{th} column is indexed by (j,k) for some k.

The matrix A has the *P-property* if the determinants of all of its representative submatrices have the same nonzero sign. (See Example 1 in Section 3 of this article.) Cottle and Dantzig [2], Theorem 3, proved the existence part of the following. (Uniqueness was proved by Szanc [13].)

Proposition 1 *Suppose that A has the P-property, and that $n_j \geq 2$ for $j = 1,2\ldots,m$. If $c \in \mathbb{R}^n$ then there exists a unique vector $v \in \mathbb{R}^m$ so that $c^T - v^T A \geq 0$, and for each $j = 1,2,\ldots,m$, $(c^T - v^T A)_{jk} = 0$ for at least one $k \in \{1,2,\ldots,n_j\}$.*

The original paper of Cottle and Dantzig [2] used a formulation with a distinguished representative submatrix. Suppose v, c and A are as in Proposition 1. Let $\hat{C}$ be the representative matrix of A for which column j is column (j, n_j) of A, for $j = 1,2,\ldots,m$. Let $\hat{A} = \hat{C}^{-1}A$ and form the matrix $\hat{A}_N$ from the columns (j,k) of $\hat{A}$ for which $k < n_j$. Let c_N be the subvector of c with subscripts (j,k) only for $k < n_j$, and let c_B be the subvector of c with subscripts (j, n_j). Then the vector $z = c_B - \hat{C}^T v$ satisfies $z \geq 0$ and $w = c_N - \hat{A}_N^T c_B + \hat{A}_N^T z \geq 0$, $z_j \prod_{k=1}^{n_j-1} w_{jk} = 0$ for $j = 1,2,\ldots,m$. Thus z is a solution to the generalized linear complementarity problem defined by matrix $\hat{A}_N^T$ and vector $c_N - \hat{A}_N^T c_B$. If $n_j = 2$ for all j, then this generalized linear complementarity problem is a standard linear complementarity problem. The columns of $\hat{A}$ not in $\hat{A}_N$ form an $m \times m$ identity matrix. This, together with the P-property of A, proves that $\hat{A}_N^T$ is a *generalized P-matrix*, i.e. each principal minor of every representative submatrix of $\hat{A}_N$ is positive.

No polynomial time algorithm is known to solve the linear complementarity problem with a P-matrix, but Megiddo [9] shows that if this problem is NP-hard then $NP = Co - NP$. His proof is easily seen to apply to the generalized LCP with a generalized P-matrix. Thus the complexity of finding the v as in Proposition 1 is unknown for general A with the P-property.

A reformulation of the P-property that follows easily from [2], Theorem 5, is that for every nonzero vector x in the row space of A, there is a $j \in \{1,2,\ldots,m\}$ so that the coordinates x_{jk} have the same nonzero sign for $k = 1,2,\ldots,n_j$. We call this the sign-preserving property of the vector x. In addition, Theorem 6 of [2] says that if A has the P-property, then the row space of A contains a positive vector.

A polytope $\mathcal{P}_{A,b}$ defined by the system $Ax = b$, $x \geq 0$ is said to be combinatorially equivalent to a product of simplices if the representative submatrices of A are the feasible bases of the system and all of the corresponding basic feasible solutions are nondegenerate. By nondegeneracy we mean that the solution to $Cx = b$ is positive for every representative submatrix C of A. If $n_j = 2$ for $j = 1,2,\ldots,m$, then such a polytope $\mathcal{P}_{A,b}$ is combinatorially equivalent to a cube. If $\mathcal{P}_{A,b}$ is combinatorially equivalent to a product of simplices, then A has the P-property. Linear programs with feasible regions combinatorially equivalent to products of simplices are prominent in the paper of Amenta and Ziegler [1].

A square matrix with nonpositive off-diagonal entries is called a *Z-matrix.* The following classical theorem [4] is central for Z-matrices.

Proposition 2 *Let M be a Z-matrix. The following are equivalent.*

1. *M is a P-matrix, i.e. all principal minors of M are positive.*

2. *The system $x > 0, Mx > 0$ has a solution.*

3. *The inverse of M exists and is nonnegative.*

A Z-matrix M satisfying the conditions of Proposition 2 is called a *K-matrix.* See, e.g., [3], Theorem 3.11.10.

Theorem 1 *Suppose A has the P-property. Then there exists an $m \times m$ matrix $\overline{X}$ so that*

1. *$(\overline{X}A)_{ijk} > 0$ whenever $i = j$.*

2. *$(\overline{X}A)_{ijk} \leq 0$ whenever $i \neq j$.*

3. *For every pair (i, j) with $i \neq j$, there exists $k \in \{1, 2, \dots, n_j\}$ so that $(\overline{X}A)_{ijk} = 0$.*

$\overline{X}$ is unique up to positive scaling of its rows.

Proof. We recall the matrix $\hat{A} = \hat{C}^{-1}A$, where $\hat{C}$ is the representative matrix of A for which column j is column (j, n_j) of A, for $j = 1, 2, \dots, m$. Let $\hat{A}^m$ be the matrix obtained from $\hat{A}$ by deleting row m and deleting columns $(m, 1), (m, 2), \dots, (m, n_m)$. Also, define c to be -1 times the vector obtained from row m of $\hat{A}$ by deleting entries $(m, 1), (m, 2), \dots, (m, n_m)$.

$$\hat{A} = \left[\begin{array}{c|c|c} \hat{A}^m & \cdots & 0 \\ \hline -c^T & \cdots & 1 \end{array}\right]$$

Then $\hat{A}^m$ has the P-property (each of its representative submatrices can be extended to a representative matrix of $\hat{A}$ with the same determinant), so there is a unique vector $v \in \mathbb{R}^{m-1}$ satisfying $c^T - v^T\hat{A}^m \geq 0$, and for every $j \neq m$ there is a $k \in \{1, 2, \dots, n_j\}$ so that $(c^T - v^T\hat{A}^m)_{jk} = 0$. Appending a 1 to the end of vector v, we get $(v^T, 1)\hat{A}$ which has the negative of the sign pattern of $c^T - v^T\hat{A}^m$ on coordinates not indexed by $(m, 1), (m, 2), \dots, (m, n_m)$, and which is 1 on component (m, n_m). The sign-preserving property of the row space of A then implies that $(v^T, 1)\hat{A}$ is positive on all components indexed by $(m, 1), (m, 2), \dots, (m, n_m)$. Thus $(v^T, 1)\hat{C}^{-1}$ can serve as row m of the matrix $\overline{X}$. A similar construction can be carried out for each of the rows of $\overline{X}$.

Suppose that $w \in \mathbb{R}^{m-1}$ and $\omega \in \mathbb{R}$ are such that $(w^T, \omega)\hat{A}$ is positive on components in block m and nonpositive on all other blocks, with at least one zero entry in each block other than block m. Then ω would have to be positive and thus we would have $c^T - \frac{1}{\omega}w^T\hat{A}^m \geq 0$, and for every $j \neq m$ a $k \in \{1, 2, \dots, n_j\}$ so that $(c^T - \frac{1}{\omega}w^T\hat{A}^m)_{jk} = 0$. By the uniqueness of v, $v = \frac{1}{\omega}w$. Thus $(w^T, \omega)\hat{A} = \omega(v^T, 1)\hat{A}$. ∎

The decomposition of $\hat{A}$ shows how to prove inductively that the assumption $n_j \geq 2$ for $j = 1, 2 \dots, m$ is not needed in Proposition 1. Suppose that $d \in \mathbb{R}^n$, and suppose $n_m = 1$.

Let $\overline{d}$ be d with coordinate $(m, 1)$ removed. If $(\overline{d} + d_{m1}c)^T - v^T \hat{A}^m$ is nonnegative, with at least one entry equal to zero in each block, then the same can be said for $d^T - (v^T, d_{m1})\hat{A}$.

Theorem 1 was proved in [10] for the case $|n_j| = 2$ for all j. The algorithm in the Proof to construct $\overline{X}$ uses Proposition 1 applied to matrices with $m-1$ rows. This improves slightly upon the original proof which applied Proposition 1 to matrices with m rows. We will call the matrix $\overline{A} = \overline{X}A$ given by the previous theorem the *complementary Z-form* of A.

Theorem 2 *Let A have the P-property. The following are equivalent:*

1. *There exists a positive vector $p \in \mathbb{R}^m$ so that $p^T\overline{A} > 0$.*
2. *There exists a matrix $X \in \mathbb{R}^{m \times m}$ and a positive vector $p \in \mathbb{R}^m$ so that XA satisfies conditions (1) and (2) of Theorem 1, and $p^T XA > 0$.*
3. *There exists a vector $b \in \mathbb{R}^m$ so that for every representative submatrix C of A the solution to $Cx = b$ is positive.*

Proof. Clearly, (1) implies (2). If (2) holds, then every representative submatrix of XA is a K-matrix by part 2 of Proposition 2. Let q be any positive vector. Then for any representative submatrix C of XA, the solution to $Cx = q$ is positive, by part 3 of Proposition 2. It follows that for every representative submatrix $D = X^{-1}C$ of A, the solution to $Dx = b$ is positive for $b = X^{-1}q$. Thus (3) holds. To show that (3) implies (1), suppose that C is a representative submatrix of $\overline{X}A$ for which $C_{mj} = 0$, $j = 1, 2, \ldots, m-1$. Such a representative submatrix is guaranteed by part 3 of Theorem 1. The solution to $Cx = \overline{X}b$ must be positive, and C_{mm} is positive, so component m of $\overline{X}b$ is positive. In a similar way, we see that all components of $\overline{X}b$ are positive. It follows that for an arbitrary representative matrix C of $\overline{A}$, we have an $x > 0$ such that $Cx = \overline{X}b > 0$, implying that C^T is a K-matrix by part 2 of Proposition 2. Part 1 of Proposition 2 shows that the transpose of a K-matrix is a K-matrix, so each representative matrix C of $\overline{A}$ is a K-matrix. We add a column $(j, n_j + 1)$, equal to the j^{th} standard basis vector of $\mathbb{R}^m$, to $\overline{A}$ for each $j = 1, 2, \ldots, m$. The resulting matrix retains the P-property, by property 1 of Proposition 2, and so it has a positive vector in its row space. There is therefore a positive vector p so that $p^T\overline{A} > 0$. ∎

The equivalence of (2) and (3) in Theorem 2 was already discussed in [6]. Condition (2) is equivalent (see [6]) to the condition that the matrix $\hat{A}_N^T$, from the discussion after Proposition 1, is a *generalized hidden K-matrix.* The equivalence of (1) to the other conditions was discussed for the case $|n_j| = 2$ for all j in [10] and [3], section 4.8. The proof above seems to be particularly straightforward.

The matrix $\overline{X}$ is useful for several reasons. If the equivalent conditions of Theorem 2 do *not* hold, $\overline{X}$ provides a quick certificate for proving this. One needs only to display $\overline{A}$ and a representative submatrix C of $\overline{A}$ for which the system $Cx \leq 0, x \geq 0, x \neq 0$ has a solution (Theorem 2.7.11 (of the Alternative) of [3]).

If X and p satisfy (2), then one can define the cone $\mathcal{C}_X := \{X^{-1}q : q > 0\}$. The proof above for (2) $\Rightarrow$ (3) shows that for every $b \in \mathcal{C}_X$, the representative submatrices of A are nondegenerate feasible bases of $Ax = b, x \geq 0$. The proof for (3) $\Rightarrow$ (1) shows that every b for which the representative submatrices of A are nondegenerate feasible bases of

$Ax = b, x \geq 0$ is in $\mathcal{C}_{\overline{X}}$. Thus $\mathcal{C}_{\overline{X}}$ is maximal among such $\mathcal{C}_X$. Linear programs to find X and p satisfying (2) if they exist have been proposed by [12] and [11], but these are not guaranteed to find $\overline{X}$, the one for which $\mathcal{C}_X$ is largest.

If A has the P-property and (3) is satisfied by some vector b, then $\overline{X}b > 0$, so the sign pattern of the matrix in the system $\overline{X}Ax = \overline{X}b$ shows that there are no feasible bases of $Ax = b, x \geq 0$ other than the representative submatrices of A. It follows that the feasible region of $Ax = b, x \geq 0$ is combinatorially equivalent to a product of simplices.

In the case where $|n_j| = 2$ for all $j = 1, 2, \ldots, m$, a vector b for which the representative submatrices of A are nondegenerate feasible bases of $Ax = b, x \geq 0$ is a so-called n-step vector for the matrix $\hat{A}_N$ (see discussion after Proposition 1) and can be used to solve any linear complementarity problem with matrix $\hat{A}_N$ in at most m pivots ([3] Section 4.8). It is therefore desirable to have a description, $\mathcal{C}_{\overline{X}}$, of the set of all such b.

The location of the zeroes in $\overline{A}$ may offer useful combinatorial information. If $|n_j| = 2$ for all j, the conditions of Theorem 2 hold, and one of the representative submatrices of $\overline{A}$ is a diagonal matrix, then the paper [5] shows that the simplex method finds the vector guaranteed by Proposition 1 very quickly.

It follows that the efficient computation of $\overline{X}$ is of interest.

3 Discounted Markov Decision Problems

A discounted Markov decision problem is defined by an $m \times n$ matrix A, with the columns indexed by pairs (j, k) as in the previous section, and a vector $c \in \mathbb{R}^n$. Every representative submatrix of A is of the form $I - \gamma P$, with I the $m \times m$ identity matrix, γ a scalar between 0 and 1 and P a nonnegative matrix with column sums equal to 1. It is easy to see that the matrix A for a discounted Markov decision problem satisfies the P-property. Moreover, it satisfies condition (2) of Theorem 2 with $X = I$ and $p = e^m$, the vector of length m with all entries equal to one. From the proof of Theorem 2, we see that for any positive vector b and every representative submatrix C of A, the solution to $Cx = b$ will be positive. It is also easy to see that there can be no other basic feasible solutions of $Ax = b$.

As stated in the introduction, the Markov decision problem is to find $v \in \mathbb{R}^m$ that maximizes $v^T b$, subject to $v^T A \leq c^T$. This problem is dual to: minimize $c^T x$ subject to $Ax = b, x \geq 0$, which is often referred to as the "primal" linear program for the Markov decision problem. Because the optimal basis to the primal linear program must be a representative submatrix of A, the complementary slackness conditions tell us that we need to find v so that $c^T - v^T A \geq 0$, and for each $j = 1, 2, \ldots, m$, $(c^T - v^T A)_{jk} = 0$ for at least one $k \in \{1, 2, \ldots, n_j\}$. We saw in section 2 that such a vector is unique when A has the P-property. It follows that the optimal basis does not depend on the vector b.

We would like to show that the complexity results for the Markov decision problem can be applied to any linear program of the form minimize $c^T x$ such that $Ax = b, x \geq 0$ if A satisfies the conditions of Theorem 2. For each representative submatrix C of the matrix A, we define v_C as the unique vector in $\mathbb{R}^m$ so that $c^T - v_C^T A$ is 0 on the components corresponding to the columns of C. The *reduced cost vector* c_C is then $c_C^T = c^T - v_C^T A$.

Proposition 3 *Suppose that the matrix A has the P-property and satisfies the equivalent conditions of Theorem 2. Let X be as in part (2) of Theorem 2. Then, for any $c \in \mathbb{R}^n$, there is a Markov decision problem, with a matrix $\widetilde{XA}$ that has one row and column more* *than A, so that for each representative matrix C of A the vector $(c_C, 0) \in \mathbb{R}^{n+1}$ is the reduced cost vector of the representative matrix of the Markov decision problem obtained by augmenting C by a row and column.*

Proof. Suppose that A satisfies the equivalent conditions of Theorem 2. Then there exists $X \in \mathbb{R}^{m\times m}$ satisfying (2) of Theorem 2, and by scaling the rows of X we can assure that $(e^n)^T \geq (e^m)^T XA > 0$. For each column (j,k), let γ_{jk} satisfy $1-\gamma_{jk} = (e^m)^T(XA)_{jk}$. Then column $(XA)_{\cdot jk}$ will equal $e_j - \gamma_{jk}p_{jk}$, where p_{jk} is a nonnegative vector with entries that sum to one. Each γ_{jk} can be interpreted as a discount factor for action (j,k). It will not generally be true that the γ_{jk} will all be equal, but we can enforce this by adding a row and column in a way that will not affect the application of any pivoting algorithm applied to the matrix XA. Let $d = min_{(j,k)}(e^m)^T(XA)_{jk}$, the smallest column sum of XA. We can augment XA to an $(m+1)\times(n+1)$ matrix $\widetilde{XA}$ by first adding an $(n+1)^{st}$ column, indexed by $(m+1,1)$, of zeroes, and then adding a row of length $n+1$ for which entry (j,k) is $d-(e^m)^T(XA)_{jk}$ for $j = 1,2,\ldots,m+1$, $k = 1,2,\ldots,n_j$.

Every representative submatrix of the matrix $\widetilde{XA}$ will contain the new column $(m+1,1)$. Suppose that c is a vector in $\mathbb{R}^n$, and let $(c^T,0)$ be c^T with a 0 appended as an $(m+1,1)$ component. Suppose that C is a representative submatrix of A with column indices in a set $J \subseteq [n]$. Let $\tilde{C}$ be the representative submatrix of $\widetilde{XA}$ with columns indexed by $J \cup \{(m+1,1)\}$. Let c_J be the subvector of c with components indexed by J, and let $(c_J,0)^T$ be the subvector of $(c^T,0)$ with components indexed by $J \cup \{(m+1,1)\}$. It is easy to see that a vector $\tilde{v} \in \mathbb{R}^{m+1}$ satisfies $\tilde{v}^T\tilde{C} = (c_J,0)^T$ if and only if $\tilde{v}^T_{m+1} = 0$ and the vector v obtained from $\tilde{v}$ by deleting the last component, satisfies $v^TXC = c_J^T$. Thus, $v^TX = v_C^T$ and we have $(c_C,0)^T = (c,0)^T - \tilde{v}^T\widetilde{XA}$, which means that each reduced cost vector for the Markov decision problem is obtained by adding a 0 component to the corresponding reduced cost vector of the original linear program. ∎

Corollary 1 *Suppose A satisfies the equivalent conditions of Theorem 2, $X \in \mathbb{R}^{n\times n}$ is as in (2) of Theorem 2, and $b = X^{-1}q$ for some positive q, and we solve the linear program minimize c^Tx such that $Ax = b, x \geq 0$ with a pivoting algorithm that bases its choice of entering variable only on the reduced cost vector of the current basis. Then the algorithm will take the same number of steps as when it is applied to a related Markov decision process with one extra row and column.*

We note that application of the simplex method with Dantzig's rule to solve the LP, minimize c^Tx such that $Ax = b, x \geq 0$, does not require us to know X, or a right hand side b. The entering variable for a given feasible basis is determined only by the reduced cost coefficients. If the entering variable is indexed by the pair (j,k), then the leaving variable is the currently basic variable indexed by (j,k') for some $k' \neq k$. One example of a linear program that Proposition 2 shows is equivalent to a Markov decision problem is the Klee-Minty cube, studied in [8]. The transformation of the Klee-Minty cube into a Markov decision problem necessarily leads to a very small value of $1-\gamma$. The paper [15] proves that Dantzig's rule of choosing the entering variable with the smallest reduced cost, applied to a discounted Markov decision problem with an $m\times n$ matrix and discount factor γ, takes at most $\frac{m(n-m)}{1-\gamma}\cdot log(\frac{m^2}{1-\gamma})$ pivots to find the solution. This bound is a polynomial in n and m if γ is not assumed to be part of the input data.

Given a matrix A with the P-property, it seems desirable to look for an equivalent Markov decision problem matrix with $1-\gamma$ as large as possible. This leads to the following linear program:

$$\begin{array}{c} \text{maximize } d \text{ subject to} \\ (XA)_{ijk} \leq 1 \text{ whenever } i = j \\ (XA)_{ijk} \leq 0 \text{ whenever } i \neq j \\ (e^m)^T XA \geq d(e^n)^T \end{array}.$$

This linear program is similar to one from [12].

We will call this linear program $LP(A)$. It is clear that $LP(A)$ always has the feasible solution $X = 0$, $d = 0$, and that $LP(A)$ has a solution with positive optimal value if and only if one of the equivalent conditions of Theorem 2 is satisfied. If $LP(A)$ has a solution with positive optimal value d, then $\gamma = 1 - d$ is the discount factor of the Markov decision problem with matrix $\widetilde{XA}$.

Example 1 *Let*

$$A = \left[\begin{array}{cc|cc|cc} 4 & 4 & -1 & -3 & -2 & -1 \\ -2 & -1 & 4 & 4 & -1 & -1 \\ -1 & -2 & -1 & 0 & 4 & 4 \end{array}\right].$$

Using Maple 11 to solve $LP(A)$ yields the solution

$$X = \left[\begin{array}{ccc} 111/350 & 37/350 & 37/350 \\ 1/25 & 7/25 & 2/25 \\ 4/35 & 3/35 & 2/7 \end{array}\right], d = 33/70,$$

and

$$XA = \left[\begin{array}{cc|cc|cc} 333/350 & 333/350 & 0 & -37/70 & -111/350 & 0 \\ -12/25 & -7/25 & 1 & 1 & -1/25 & 0 \\ 0 & -1/5 & -2/35 & 0 & 29/35 & 33/35 \end{array}\right].$$

In this example, the columns appear in the order $(1,1), (1,2), (2,1), (2,2), (3,1), (3,2)$, and $n_1 = n_2 = n_3 = 2$. Note that the matrix A already satisfies the sign conditions (1) and (2) of Theorem 1, and it satisfies condition (2) of Theorem 2. The matrix XA found by Maple 11 does not, however, satisfy condition (3) of Theorem 1, because neither $(XA)_{211}$ nor $(XA)_{212}$ is 0. This leads one to ask if there must be an optimum solution of $LP(A)$, if the objective value of $LP(A)$ is positive, for which all three conditions of Theorem 1 are satisfied.

4 The two step method

We will assume in this section that A has the P-property and that A satisfies the equivalent conditions of Theorem 2. We want to show that $LP(A)$ can be computed in strongly polynomial time if the optimal value of $LP(A)$ is a constant and not part of the input. In addition, an optimal X satisfying all three conditions of Theorem 1 will be found. First we recall the matrix $\hat{A}^m$ from the proof of Theorem 1.

Lemma 1 *The matrix $\hat{A}^m$ satisfies the equivalent conditions of Theorem 2.*

Proof. Suppose $\overline{X}$ is as in Theorem 1. Element $(\overline{X}A)_{mmn_m}$ is positive. Let E be the matrix obtained from the $m \times m$ identity matrix by replacing entry (j, m) by $-(\overline{X}A)_{jmn_m}/(\overline{X}A)_{mmn_m}$ for $j = 1, 2, \ldots, m-1$. Then the matrix $E\overline{X}A$ has, just as $\hat{A}$ does, as its last column a multiple of e_m, column m of the $m \times m$ identity matrix. It is easy to see that if $p > 0$ and $p^T\overline{X}A > 0$, then p^TE^{-1} will be positive and satisfy $(p^TE^{-1})(E\overline{X}A) > 0$. Form $\overline{A}^m$ from $E\overline{X}A$ by deleting row m and column $(m, 1), (m, 2), \ldots, (m, n_m)$. Matrix $\overline{A}^m$ satisfies the equivalent conditions of Theorem 2. The row spaces of $\hat{A}^m$ and $\overline{A}^m$ are the same, so there exists an $(m-1) \times (m-1)$ matrix X^m so that $\overline{A}^m = X^m\hat{A}^m$. ∎

Proposition 4 *Row m of $\overline{X}$ can be found by solving a linear program that is equivalent to solving a Markov decision problem.*

Proof. As in the proof of Theorem 1, define c to be -1 times the vector obtained from row m of $\hat{A}$ by deleting entries $(m, 1), (m, 2), \ldots, (m, n_m)$. Because $\hat{A}^m$ satisfies the equivalent conditions of Theorem 2, there exists a vector $b \in \mathbb{R}^{m-1}$ so that the representative submatrices of $\hat{A}^m$ are the nondegenerate feasible bases of $\hat{A}^m x = b, x \geq 0$. The vector v that was found in the proof of Theorem 1 can be found as the dual solution to the linear program: minimize c^Tx subject to $\hat{A}^m x = b, x \geq 0$. Recall that one does not need to know b in order to solve this linear program by a pivoting method. For $(v, 1)$ obtained by appending a 1 to v we obtain $(v, 1)\hat{A}$, a vector in the row space of A that is nonpositive on entries not indexed by $(m, 1), (m, 2), \ldots, (m, n_m)$ and has for each $j = 1, 2, \ldots, m-1$ an entry (j, k) that is zero. It is also positive on entry (m, n_m), so by the sign-preserving property of the row space of A it is positive on all components $(m, 1), (m, 2), \ldots, (m, n_m)$. Thus $(v, 1)^T\hat{C}^{-1}$ can serve as row m of matrix $\overline{X}$. ∎

We can obtain the other rows of $\overline{X}$ similarly, by solving linear programs.

Example 2 *Applied to the previous example, we obtain*

$$\overline{X} = \begin{bmatrix} 1/3 & 1/9 & 1/9 \\ 3/19 & 7/19 & 5/38 \\ 4/33 & 1/11 & 10/33 \end{bmatrix}$$

and

$$\overline{X}A = \left[\begin{array}{cc|cc|cc} 1 & 1 & 0 & -5/9 & -1/3 & 0 \\ -9/38 & 0 & 45/38 & 1 & -3/19 & 0 \\ 0 & -7/33 & -2/33 & 0 & 29/33 & 1 \end{array}\right].$$

Note that the entry in position (j, j, n_j) is 1 for all j.

The second step of our method is to determine an optimum scaling of the rows of the matrix $\overline{X}$. Let $\overline{A} = \overline{X}A$. The column sums of $\overline{A}$ are not necessarily between 0 and 1. We find an optimum scaling vector $x \in \mathbb{R}^m$ by solving the linear program:

$$\begin{array}{c} \text{maximize } d \text{ subject to} \\ x^T\overline{A} \geq d(e^n)^T \\ x_j\overline{A}_{jjk} \leq 1 \text{ for all } j = 1, 2, \ldots, m, k \in [n_j]. \end{array}$$

If $diag(x)$ is the $m \times m$ diagonal matrix with x on the diagonal, we see that $x^T\overline{A} = (e^m)^T diag(x)\overline{A}$, so the vector x can be thought of as an optimal scaling of the rows of $\overline{A}$

to bring it into the MDP format. We call this linear program the scaling LP. The matrix $diag(x)$ is feasible for $LP(\overline{A})$, and it has the extra feature that condition (3) of Theorem 1 is satisfied for matrix $diag(x)\overline{A}$. It is not immediately obvious, though, that the optimal value of the scaling LP is equal to the optimal value of $LP(A)$. To prove the equality is our remaining goal.

Example 3 *Applying the scaling LP to $\overline{A}$ from our previous example yields*

$$diag(x) = \begin{bmatrix} 47/70 & 0 & 0 \\ 0 & 38/45 & 0 \\ 0 & 0 & 33/35 \end{bmatrix}, d = 33/70$$

and

$$diag(x)\overline{X}A = \left[\begin{array}{cc|cc|cc} 47/70 & 47/70 & 0 & -47/126 & -47/210 & 0 \\ -1/5 & 0 & 1 & 38/45 & -2/15 & 0 \\ 0 & -1/5 & -2/35 & 0 & 29/35 & 33/35 \end{array}\right].$$

Proposition 5 *The linear program $LP(A)$ has the same optimal value as the scaling LP.*

Proof. It is clear that the linear programs $LP(A)$ and $LP(\overline{A})$ have the same optimal value, so we want to show that the optimal values of $LP(\overline{A})$ and the scaling LP are the same. Note that the optimal values of both LPs are positive. For the scaling LP, this implies, by (3) of Proposition 2, that the solution x will be nonnegative. We will compare the dual linear programs. The dual to $LP(\overline{A})$ has variables y_{ijk} corresponding to entries of $\overline{A}$ and variables w_{jk} corresponding to columns of $\overline{A}$.

$$\begin{array}{c} \text{minimize } \sum_{i=j\in[m],k\in[n_j]} y_{ijk} \text{ subject to} \\ \sum_{j\in[m],k\in[n_j]} \overline{A}_{i'jk}y_{ijk} = \sum_{j\in[m],k\in[n_j]} \overline{A}_{i'jk}w_{jk} \text{ for all } i \in [m], i' \in [m] \\ \sum_{j\in[m],k\in[n_j]} w_{jk} = 1, \\ y \geq 0, w \geq 0 \end{array}.$$

The scaling LP has variables y_{jjk} corresponding to entries of $\overline{A}$ for which the first two indices are equal, and variables w_{jk} corresponding to columns of $\overline{A}$.

$$\begin{array}{c} \text{minimize } \sum_{i=j\in[m],k\in[n_j]} y_{ijk} \text{ subject to} \\ \sum_{k\in[n_i]} \overline{A}_{iik}y_{iik} = \sum_{j\in[m],k\in[n_j]} \overline{A}_{ijk}w_{jk} \text{ for all } i \in [m] \\ \sum_{j\in[m],k\in[n_j]} w_{jk} = 1, \\ y \geq 0, w \geq 0 \end{array}.$$

Suppose that $\{y^*_{ijk} : i = j\}$ and $\{w^*_{jk}\}$ are an optimal solution to the dual of the scaling LP. We will create a solution $(\overline{y}, \overline{w})$ to the dual of $LP(\overline{A})$ by solving systems of equations. We will keep the solution to the scaling LP: $\overline{y}_{ijk} = y^*_{ijk}$ whenever $i = j$ and $\overline{w} = w^*$.

For each $i = 1, 2, \ldots, m$, we will form a linear system to solve for variables $\overline{y}_{ijk}$ for $j \neq i$. Given i, let K_i be a function that assigns to each $j \neq i$ a pair (j, k) for which $\overline{A}_{ijk} = 0$. System i will have an equation for each $i' \neq i$:

$$\sum_{j\neq i} \overline{A}_{i'jK_i(j)}\overline{y}_{ijK_i(j)} = \sum_{j\in[m],k\in[n_j]} \overline{A}_{i'jk}w^*_{jk} - \sum_{k\in[n_i]} \overline{A}_{i'ik}y^*_{iik},$$

It is easy to verify that the right hand side of this system is nonnegative. The coefficient matrix of the system is an $(m-1)\times(m-1)$ K-matrix, so the system will have a nonnegative solution when the right hand side is nonnegative. The systems for different i have disjoint sets of variables, so they are solved independently. All variables $\overline{y}_{ijk}$ not specified by the scaling LP or by the systems are set to 0. The resulting $(\overline{y}, \overline{w})$ is a feasible solution to the dual LP of $LP(\overline{A})$. The objective value of this dual solution is the same as that of the scaling LP, because the variables in the two objective functions are the same. ∎

5 Complexity

Now we assume that $(\overline{X}, d^*)$ solves $LP(A)$ and that $d^* > 0$. We will assume that the matrix $\overline{X}$ satisfies conditions (1) – (3) of Theorem 1. We have $(e^m)^T\overline{X}A \geq d^*(e^n)^T$.

Proposition 6 *Row m of $\overline{X}$ is found by a strongly polynomial algorithm, solving a linear program equivalent to a Markov decision problem with an $m \times n$ matrix, if d^* is considered to be a constant.*

Proof. We need to show that the discount factor for the Markov decision problem equivalent to the linear program of the proof of Proposition 4 is no larger than $1 - d^*$. We prove this using the matrix $\overline{A}^m$ from Lemma 1, which is row equivalent to $\hat{A}^m$. Then for any (j, k) with $j \in [m-1]$ and $k \in [n_j]$, we have $d^* \leq (e^m)^T\overline{A}_{\cdot jk} = (e^{m-1})^T\overline{A}^m_{\cdot jk} + (1 + \sum_{i=1}^{m-1}(\overline{X}A)_{imn_m}/(\overline{X}A)_{mmn_m})(\overline{X}A)_{mjk}$. This last term is nonpositive, so $(e^{m-1})^T\overline{A}^m_{\cdot jk} \geq d^*$. Recall that we defined X^m by $\overline{A}^m = X^m\hat{A}^m$. While the matrix $\overline{A}^m$ is not necessarily the complementary Z-form of $\hat{A}$, the pair (X^m, d^*) is a feasible solution to $LP(\hat{A}^m)$. Therefore, the optimal value of $LP(\hat{A}^m)$ is at least d^*. It follows from [15] that the linear program to compute row m of $\overline{X}$ (without scaling) can be solved in at most $\frac{m(n-m)}{d^*} \cdot log(\frac{m^2}{d^*})$ pivots using Dantzig's rule. ∎

Proposition 7 *The scaling LP can be solved by applying the strongly polynomial algorithm of [15], again assuming d is a constant, followed by solving a one-variable linear program.*

Proof. Here we first find a vector v satisfying $v^T\overline{A} \geq (e^n)^T$, such that for every $j \in [m]$ there exists a $k \in [n_j]$ such that component (j, k) of $(e^n)^T - v^T\overline{A}$ is zero. Note that $v^T\overline{A} \geq (e^n)^T$ is equivalent to $(-v^T)\overline{A} \leq (-e^n)^T$, so finding $-v^T$ is, as before, equivalent to solving a Markov decision problem with an $m \times n$ matrix and discount factor $1 - d^*$. Then we find $\hat{d}$ which maximizes d for which $dv_i\overline{A}_{iik} \leq 1$ for all $i \in [m]$, $k \in [n_i]$. Let (i^*, i^*, k^*) index an entry of $\hat{d}(diag(v))\overline{A}$ that is equal to 1. Let C be a representative submatrix of $\hat{d}(diag(v))\overline{A}$ satisfying $(e^m)^TC = \hat{d}(e^m)^T$. We build a dual solution (y, w) to the scaling LP as follows. We let $y_{i^*i^*k^*} = \hat{d}$. Let $\hat{w}$ solve $C\hat{w} = \hat{d}e_{i^*}$ and let the components of w corresponding to the columns of C equal the corresponding components of $\hat{w}$. Let all other components of y and w be 0. Then the system $C\hat{w} = \hat{d}e_{i^*}$ guarantees that the constraints

$$\sum_{k\in[n_i]} \overline{A}_{iik}y_{iik} = \sum_{j\in[m],k\in[n_j]} \overline{A}_{ijk}w_{jk} \text{ for all } i \in [m]$$

will be satisfied. For example, in equation i^* we have $\sum_{k\in[n_{i^*}]}\overline{A}_{i^*i^*k}y_{i^*i^*k} = \frac{1}{\hat{d}v_{i^*}}(\hat{d}v_{i^*}\overline{A}_{i^*i^*k^*})\hat{d} = \frac{1}{\hat{d}v_{i^*}}\hat{d} = \frac{1}{\hat{d}v_{i^*}}(C\hat{w})_{i^*} = \sum_{j\in[m],k\in[n_j]}\overline{A}_{i^*jk}w_{jk}$. Premultiplying both sides of $C\hat{w} = \hat{d}e_{i^*}$ by $(e^m)^T$ gives $(e^m)^TC\hat{w} = \hat{d}(e^m)^Te_{i^*}$, or $(e^m)^T\hat{w} = 1$. Thus (y, w) is a feasible solution to the dual of the scaling LP, which has the same value as the solution $\hat{d}v$ of the primal LP. ∎

Theorem 3 *If the optimal value of the linear program $LP(A)$ is positive, and this optimal value is considered a constant that is not part of the input data, then there is a strongly polynomial algorithm to solve $LP(A)$.*

It is interesting that the two step method described here serves two seemingly independent purposes, to satisfy condition (3) of Theorem 1 and to be efficient. There does not appear to be reason to assume that the straightforward application of, say, the simplex method to $LP(A)$ would be strongly polynomial. The solution $diag(x)\overline{X}$ found to $LP(A)$ also has the property that for each $j \in [m]$ there is an index (j, k) for which $(x^T\overline{X}A)_{jk} = d$, where d is the optimal value of the LP.

6 Extensions

For a matrix A with the P-property and a vector $c \in \mathbb{R}^n$, the complexity of finding the vector v satisfying $(c^T - v^TA)_{jk} \geq 0$ for each $j \in [m]$, $(c^T - v^TA)_{jk} = 0$ for some $k \in \{1, 2, \ldots, n_j\}$ is drastically better for A satisfying the equivalent conditions of Theorem 2 than it is for general A with the P-property. It should be pointed out that Ye [14] had already in 2005 shown that the Markov decision problem can be solved in $O(n^4log(\frac{1}{1-\gamma}))$ arithmetic operations, using an interior point method. Note that the dependence on $\frac{1}{1-\gamma}$ is logarithmic here, rather than linear.

The unresolved complexity of the (generalized) linear complementarity problem with a (generalized) P-matrix makes one eager to escape the confines of Theorem 2.

The paper [7] solves a considerable generalization of the Markov decision problem with complexity similar to that of [15]. Here one assumes that A and c are as in a Markov decision problem. The generalization allows for a partition of the set $[m]$ into two parts M_1 and M_2, and the vector v must satisfy $(c^T - v^TA)_{jk} \geq 0$ for $j \in M_1$ and $(c^T - v^TA)_{jk} \leq 0$ for $j \in M_2$, in addition to the complementarity condition which says that for each $j \in [m]$, $(c^T - v^TA)_{jk} = 0$ for some $k \in \{1, 2, \ldots, n_j\}$. No linear programming formulation of this problem is known. Define the matrix $A_{\overline{M_2}}$ and vector $c^T_{\overline{M_2}}$ to be obtained from A and c^T by negating the columns and entries in blocks indexed by M_2. Then the v solving the generalization satisfies $c^T_{\overline{M_2}} - v^TA_{\overline{M_2}} \geq 0$ in addition to the complementarity constraints. That is, v solves the generalized linear complementarity problem. The matrix $A_{\overline{M_2}}$ does not necessarily satisfy the conditions of Proposition 2.

Call a matrix A with the P property a *game-type* matrix if there exists a partition of $[m]$ into two parts M_1 and M_2 so that the matrix $A_{\overline{M_2}}$ satisfies the equivalent conditions of Theorem 2. Equivalently, the optimum value d of $LP(A_{\overline{M_2}})$ is positive. An interesting open problem is to find a strongly polynomial time algorithm that takes a game-type matrix A and $c \in \mathbb{R}^n$, and finds v as in Proposition 1 assuming the partition (M_1, M_2) is not known but assuming d is a fixed positive constant.

We say the matrix A has the *diagonal dominance* property if A satisfies $1 \geq A_{jjk} > \sum_{i \neq j} |A_{ijk}|$ for every column (j,k). It is possible to show that the diagonal dominance property implies the P-property (See [11] for a simple proof.) and that for every game-type matrix A there is a square matrix X so that XA has the diagonal dominance property. If A has the diagonal dominance property, let d satisfy $A_{jjk} - \sum_{i \neq j} |A_{ijk}| \geq d$ for all (j,k). A more general open problem is to find, given A and $c \in \mathbb{R}^n$, a strongly polynomial time algorithm to find v as in Proposition 1 assuming d is a fixed positive constant.

There exist A with the P-property for which there is no square X so that XA has the diagonal dominance property (see [11]), so solving these open problems will not find a v as in Proposition 1 efficiently for all matrices with the P-property.

7 Acknowledgement

The author would like to thank the referees for their helpful comments.

References

[1] N. Amenta, G. Ziegler, "Deformed products and maximal shadows of polytopes" *Advances in Discrete and Computational Geometry,* Amer. Math. Soc., Providence, Contemprary Mathematics **223**, (1996), 57–90.

[2] R.W. Cottle, G.B. Dantzig, "A generalization of the linear complementarity problem," *Journal of Combinatorial Theory* **8** (1970), 79–90.

[3] R.W. Cottle, J.-S. Pang, R. E. Stone, "The Linear Complementarity Problem," Academic Press, Boston, 1992.

[4] M. Fiedler, V. Ptak, "On matrices with non-positive off-diagonal elements and positive principal minors," *Czechoslovak Mathematical Journal*, **12** (1962), 382–400.

[5] J. Foniok, K. Fukuda, B. Gärtner, H.-J. Lüthi, "Pivoting in linear complementarity: two polynomial-time cases," Discrete and Computational Geometry **42** (2009), 187–205

[6] B. Gärtner, W.D. Morris, Jr., L. Rüst, "Unique sink orientations of grids," *Algorithmica* **51** (2008), 200–235.

[7] T.D. Hansen, P.B. Miltersen, U. Zwick, "Strategy iteration is strongly polynomial for 2-player turn-based stochastic games with a constant discount factor," *Journal of the ACM* **60** (2013)

[8] T. Kitahara, S. Mizuno "Klee-Minty's LP and upper bounds for Dantzig's simplex method," *Operations Research Letters* **39** (2011), 88–91.

[9] N. Megiddo, "A Note on the Complexity of P Matrix LCP and Computing an Equilibrium," Research Report RJ 6439, IBM Almaden Research Center, San Jose, California, 1988.

[10] W.D. Morris, Jr., J. Lawrence, "Geometric properties of hidden Minkowski matrices," *SIAM Journal on Matrix Analysis and Applications* **10** (1989), 229–232.

[11] W.D. Morris, Jr., M. Namiki, "Good hidden P-matrix sandwiches," *Linear Algebra and its Applications* **426** (2007), 325–341.

[12] J.-S. Pang, "On discovering hidden Z-matrices," in (C.V. Coffman, G.J. Fix, eds.) *Constructive approaches to Mathematical Models,* Academic Press, New York, 1979, 231–241.

[13] B.P. Szanc, "The generalized complementarity problem," Ph. D. Thesis, Rensselaer Polytechnic Institute, Troy, NY, 1989.

[14] Y. Ye, "A new complexity result on solving the Markov decision problem," *Mathematics of Operations Research* **30** (2005), 733–749.

[15] Y. Ye, "The simplex and policy-iteration methods are strongly polynomial for the Markov decision problem with a fixed discount rate" *Mathematics of Operations Research* **4** (2011), 593–603.