Electronic Journal of Linear Algebra ISSN 1081-3810
A publication of the International Linear Algebra Society
Volume 13, pp. 131-145, April 2005

ELA
www.math.technion.ac.il/iic/ela

A PARALLEL ALGORITHM FOR COMPUTING THE GROUP INVERSE VIA PERRON COMPLEMENTATION[*]

MICHAEL NEUMANN[†] AND JIANHONG XU[‡]

Abstract. A parallel algorithm is presented for computing the group inverse of a singular M–matrix of the form $A = I - T$, where $T \in \mathbb{R}^{n\times n}$ is irreducible and stochastic. The algorithm is constructed in the spirit of Meyer's Perron complementation approach to computing the Perron vector of an irreducible nonnegative matrix. The asymptotic number of multiplication operations that is necessary to implement the algorithm is analyzed, which shows that the algorithm saves a significant amount of computation over the direct computation of the group inverse of A.

Key words. Group inverses, M–matrices, Perron complements, Parallel algorithm.

AMS subject classifications. 15A09, 15A48, 15A51, 65F30, 65C40.

1. Introduction. The *group inverse* of a matrix $A \in \mathbb{R}^{n\times n}$, denoted by $A^{\#}$ when it exists, is the unique matrix $X \in \mathbb{R}^{n\times n}$ that satisfies the matrix equations $AXA = A$, $XAX = X$, and $AX = XA$. The group inverse has been utilized in a variety of applications [5, 6, 7, 8, 9, 12, 14, 18, 19, 20, 21, 22, 23, 25, 26, 30], mostly in the context of a singular M–matrix

$$A = \rho_S I - S, \tag{1.1}$$

where $S \in \mathbb{R}^{n\times n}$ is nonnegative and irreducible and where ρ_S is its Perron root.[1] In particular, these applications include:

(i) For a finite ergodic Markov chain with a transition matrix T, Meyer [25] has shown that virtually all the important characteristics of the chain can be determined from the group inverse of $A = \rho_T I - T = I - T$. Furthermore, in [8, 14, 21, 26, 30] it has been shown that the entries of $A^{\#}$ can be used to provide perturbation bounds on the stationary distribution vector of the chain.

(ii) In [5, 6] it has been shown that for A in (1.1), the signs of the entries of $A^{\#}$ give us qualitative information about the behavior of the Perron root of S as a function of the entries of S, namely in addition to it being known that the Perron root is a strictly increasing function of each of the entries, the signs of the entries of $A^{\#}$ answer the question whether the Perron root is a concave or a convex function of each of the entries. Results of this type have been applied in the study of matrix population models in [19, 23].

[*]Received by the editors 20 December 2004. Accepted for publication 5 April 2005. Handling Editor: Ravindra B. Bapat.

[†]Department of Mathematics, University of Connecticut, Storrs, Connecticut 06269–3009, USA (neumann@math.uconn.edu). The work of this author was supported in part by NSF Grant No. DMS0201333.

[‡]Department of Mathematics and Statistics, University of West Florida, Pensacola, Florida 32514, USA (jxu@uwf.edu).

[1]For such a matrix A, the existence of $A^{\#}$ is guaranteed by the structure of the Jordan blocks of S [4, 25].

Electronic Journal of Linear Algebra ISSN 1081-3810
A publication of the International Linear Algebra Society
Volume 13, pp. 131-145, April 2005

In view of the above–mentioned applications, a question of interest is how to efficiently compute the group inverse $A^{\#}$ of A in (1.1). Several algorithms for computing $A^{\#}$ have been suggested in [1, 2, 4, 8, 11, 13, 25]. These methods typically require $O(n^3)$ to $O(n^4)$ arithmetic operations and so they can be quite expensive to implement. However, the issue of the parallel computation of $A^{\#}$ has received little attention in the literature.

In this paper we shall present an algorithm for computing $A^{\#}$ in parallel, assuming that A is given by (1.1). The numerical evidence which we shall provide indicates that this algorithm is more efficient compared with direct computation of $A^{\#}$. From the examples that we have tested with MATLAB, the savings in the number of flops (floating point operations) are roughly 50% for matrices of size ranging from $n = 50$ to $n = 1600$. An asymptotic analysis shows that the algorithm saves approximately 12.5% of the multiplication operations if it is implemented in a purely serial fashion.

A useful observation is that A in (1.1) can be reduced, using a diagonal similarity transformation (see, for example, [3, Theorem 2.5.4]), to the case where $A := I - T$, with $T \in \mathbb{R}^{n\times n}$ being irreducible and stochastic. Consequently, without loss of generality, we shall focus on this case from here on.

The key to constructing our algorithm is *Perron complementation*. Let α and β be nonempty subsets of the index set $\langle n\rangle := \{1, 2, \ldots, n\}$, both consisting of strictly increasing integers. For any $X \in \mathbb{R}^{n\times n}$, we shall denote by $X[\alpha, \beta]$ the submatrix of X whose rows and columns are determined by α and β, respectively. In the special case when $\beta = \alpha$, we shall use $X[\alpha]$ to denote $X[\alpha, \alpha]$.[2] Given the irreducible stochastic matrix $T \in \mathbb{R}^{n\times n}$, the *Perron complement* of $T[\alpha]$ is defined to be[3]

$$P_\alpha := T[\alpha] + T[\alpha, \langle n\rangle\backslash\alpha](I - T[\langle n\rangle\backslash\alpha])^{-1}T[\langle n\rangle\backslash\alpha, \alpha]. \tag{1.2}$$

Meyer [27, 28] proved that the Perron complement P_α resembles T. Specifically, P_α is also irreducible and stochastic and, in addition, on letting $\pi \in \mathbb{R}^n$ be the normalized left Perron vector of T, viz. the column vector satisfying

$$\pi^t T \;=\; \pi^t \;\text{ and }\; \|\pi\|_1 \;=\; 1, \tag{1.3}$$

then

$$\pi_\alpha \;=\; \frac{\pi[\alpha]}{\xi_\alpha}, \;\text{ with }\; \xi_\alpha \;=\; \|\pi[\alpha]\|_1, \tag{1.4}$$

where $\pi[\alpha]$ is the subvector of π determined by α, turns out to be the normalized left Perron vector of P_α. Based on these results, Meyer introduced a parallel algorithm for computing the left Perron vector via Perron complementation.

Note that the irreducible stochastic matrix T can be thought of as the transition matrix of some finite ergodic Markov chain $\{X_i\}_{i=0}^{\infty}$ on state space $\mathcal{S} = \{1, 2, \ldots, n\}$.

[2] We follow the notation in [15] for submatrices. This notation allows conclusions to be readily extended to the case when α and β consist of nonconsecutive integers.

[3] The notion of Perron complement can be similarly defined on an irreducible nonnegative matrix (see [27, 28]).

Electronic Journal of Linear Algebra ISSN 1081-3810
A publication of the International Linear Algebra Society
Volume 13, pp. 131-145, April 2005

ELA
www.math.technion.ac.il/iic/ela

The *mean first passage time* from states i to j, denoted by $m_{i,j}$, is defined to be the expected number of time steps that the chain, starting initially from state i, would take before it reaches state j for the first time [17], i.e.

$$m_{i,j} := \mathrm{E}(F_{i,j}) \;=\; \sum_{k=1}^{\infty} k\mathrm{Pr}(F_{i,j} = k), \tag{1.5}$$

where $F_{i,j} := \min\{\ell \geq 1 : X_\ell = j | X_0 = i\}$. The matrix $M = [m_{i,j}] \in \mathbb{R}^{n\times n}$ is called the *mean first passage matrix* of the chain or, simply, of T. Obviously the mean first passage matrix of the Perron complement P_α can be similarly defined since P_α can also be regarded as the transition matrix of an ergodic chain with fewer states.

As a further development of Meyer's results on Perron complementation, Kirkland, Neumann, and Xu [22] showed that the mean first passage matrix of P_α is closely related to $M[\alpha]$ and $M[\langle n\rangle\backslash\alpha]$, i.e. the corresponding submatrices of the mean first passage matrix of T. Accordingly, the mean first passage matrix of the entire chain can be computed in parallel via Perron complementation.

The interesting relationship between the irreducible stochastic matrix T and its Perron complement P_α, as shown in [22, 27, 28], leads us to explore if Perron complementation can be exploited for the parallel computation of $A^{\#}$. Specifically, on letting $B_\alpha := I - P_\alpha$, we ask the following questions: (i) How are $A^{\#}$ and $B_\alpha^{\#}$ related? (ii) Does the connection between $A^{\#}$ and $B_\alpha^{\#}$ allow the computation of $A^{\#}$ to be carried out in parallel? These questions will be completely answered in this paper. We point out that the relationship between $A^{\#}[\alpha]$ and $B_\alpha^{\#}$ has been shown in [29] for the special case when $\alpha = \{1, 2, \ldots, n-1\}$. Furthermore, for the special case when T is periodic, Kirkland [18] has developed, without recourse to the terminology of Perron complementation, formulae for the blocks of $A^{\#}$. Thus our results here will also generalize those in [18, 29].

The plan of this paper is as follows. In Section 2 we shall summarize some necessary results on Perron complementation from the literature. Our main results on the computation of the group inverse via Perron complementation will be presented in Section 3. Section 4 will be devoted to describing our parallel algorithm for computing the group inverse. Finally, some concluding remarks are given in Section 5.

2. Preliminaries. Recall that the purpose of this paper is the computation in parallel of the group inverse of $A := I - T$, where $T \in \mathbb{R}^{n\times n}$ is irreducible and stochastic.

Partition T into a $k \times k$ ($k \geq 2$) block form as follows:

$$T \;=\; \begin{bmatrix} T[\alpha_1] & T[\alpha_1, \alpha_2] & \cdots & T[\alpha_1, \alpha_k] \\ T[\alpha_2, \alpha_1] & T[\alpha_2] & \cdots & T[\alpha_2, \alpha_k] \\ \vdots & \vdots & \ddots & \vdots \\ T[\alpha_k, \alpha_1] & T[\alpha_k, \alpha_2] & \cdots & T[\alpha_k] \end{bmatrix}, \tag{2.1}$$

where $\alpha_1, \alpha_2, \ldots, \alpha_k$ are nonempty disjoint subsets of the index set $\langle n\rangle := \{1, \ldots, n\}$ which form a partition of $\langle n\rangle$. Define $W := e\pi^t$, where π is the left Perron vector of

Electronic Journal of Linear Algebra ISSN 1081-3810
A publication of the International Linear Algebra Society
Volume 13, pp. 131-145, April 2005

ELA
www.math.technion.ac.il/iic/ela

T normalized so that $\|\pi\|_1 = 1$ and where $e \in \mathbb{R}^n$ is a column vector of all ones. The vector π is the *stationary distribution vector* of the underlying Markov chain. Let $M \in \mathbb{R}^{n\times n}$ be the mean first passage matrix of T as defined in (1.5). We shall always partition $A^{\#}$, W, and M in conformity with T.

Throughout the sequel, the letters J and e will represent a matrix of all ones and a column vector of all ones, respectively, whose dimensions can be determined from the context. For any square matrix X, we shall denote by X_d the diagonal matrix obtained from X by setting all its off–diagonal entries to zero.

We begin with the following two lemmas from [4, 25].

LEMMA 2.1. ([4, Theorem 8.5.5])

$$\pi^t A^{\#} \;=\; 0. \tag{2.2}$$

LEMMA 2.2. ([25, Theorem 3.3]) *The mean first passage matrix M of T can be expressed as*

$$M \;=\; \Big[I - A^{\#} + J(A^{\#})_d\Big](W_d)^{-1}. \tag{2.3}$$

Lemma 2.2 gives the relationship between M and $A^{\#}$. For $i = 1, 2, \ldots, k$, let P_{α_i} be the Perron complement of $T[\alpha_i]$. In what follows, we shall use the subscript α_i to denote quantities associated with P_{α_i}. Put $B_{\alpha_i} := I - P_{\alpha_i}$. Since P_{α_i} is also irreducible and stochastic, its mean first passage matrix M_{α_i} is defined in the same way as M. Let π_{α_i} be the normalized left Perron vector of P_{α_i} and set $W_{\alpha_i} := e\pi^t_{\alpha_i}$. The following corollary from Lemma 2.2 characterizes the relationship between M_{α_i} and $B^{\#}_{\alpha_i}$.

COROLLARY 2.3. *For $i = 1, 2, \ldots, k$, the mean first passage matrix M_{α_i} of P_{α_i} can be expressed as*

$$M_{\alpha_i} \;=\; \Big[I - B^{\#}_{\alpha_i} + J(B^{\#}_{\alpha_i})_d\Big][(W_{\alpha_i})_d]^{-1}. \tag{2.4}$$

As already mentioned earlier, according to Meyer [27, Theorem 2.1], for $i = 1, 2, \ldots, k$, the normalized left Perron vector of P_{α_i} is given by

$$\pi_{\alpha_i} = \frac{\pi[\alpha_i]}{\xi_{\alpha_i}} \tag{2.5}$$

with $\xi_{\alpha_i} = \|\pi[\alpha_i]\|_1$. The number ξ_{α_i} is called the i–th *coupling factor*. Note that $\sum_{i=1}^k \xi_{\alpha_i} = 1$. The next lemma, again due to Meyer, points out how the coupling factors can be obtained without actually computing π.

LEMMA 2.4. ([27, Theorem 3.2]) *Let $C = [c_{i,j}] \in \mathbb{R}^{k\times k}$ be defined via*

$$c_{i,j} :=\; \pi^t_{\alpha_i} T[\alpha_i, \alpha_j]e, \quad i, j = 1, 2, \ldots, k. \tag{2.6}$$

Then C is irreducible and stochastic. Moreover, $\xi = [\xi_{\alpha_1}, \xi_{\alpha_2}, \ldots, \xi_{\alpha_k}]^t$ is the normalized left Perron vector of C.

Electronic Journal of Linear Algebra ISSN 1081-3810
A publication of the International Linear Algebra Society
Volume 13, pp. 131-145, April 2005

ELA
www.math.technion.ac.il/iic/ela

The matrix C is called the *coupling matrix.*

In the sequel we shall also require the following two lemmas from Kirkland, Neumann, and Xu [22] on the relationship between M and M_{α_i}, the mean first passage matrix of P_{α_i}. In the first lemma the principal submatrices of M are determined, while in the second lemma, the off–diagonal blocks of M are determined.

LEMMA 2.5. ([22, Theorem 2.2]) *For* $i = 1, 2, \ldots, k$,

$$M[\alpha_i] \;=\; \frac{1}{\xi_{\alpha_i}} M_{\alpha_i} + V_{\alpha_i}, \tag{2.7}$$

where

$$\begin{aligned} V_{\alpha_i} := {} & B_{\alpha_i}^{\#} T[\alpha_i, \langle n\rangle\backslash\alpha_i](I - T[\langle n\rangle\backslash\alpha_i])^{-1} J \\ & - \Big[B_{\alpha_i}^{\#} T[\alpha_i, \langle n\rangle\backslash\alpha_i](I - T[\langle n\rangle\backslash\alpha_i])^{-1} J \Big]^t. \end{aligned} \tag{2.8}$$

The matrix V_{α_i} is clearly skew–symmetric and of rank at most 2.

LEMMA 2.6. ([22, Theorem 2.3]) *For* $i = 1, 2, \ldots, k$,

$$\begin{aligned} M[\langle n\rangle\backslash\alpha_i, \alpha_i] = (I - T[\langle n\rangle\backslash\alpha_i])^{-1} \Big[& T[\langle n\rangle\backslash\alpha_i, \alpha_i] M[\alpha_i] + J \\ & - T[\langle n\rangle\backslash\alpha_i, \alpha_i](M[\alpha_i])_d \Big]. \end{aligned} \tag{2.9}$$

We comment that Corollary 2.3, formula (2.5), and Lemmas 2.4, 2.5, and 2.6 continue to hold even when any subset α_i consists of nonconsecutive indices.

3. Main Results. In this section we develop our main results in which we show how the blocks of $A^{\#}$, where $A := I - T$ with $T \in \mathbb{R}^{n\times n}$ being irreducible and stochastic, can be linked to the group inverses associated with the smaller size irreducible transition matrices arising from Perron complementation. Again we use a $k \times k$ partitioning of T which is given by (2.1).

We construct a matrix $U \in \mathbb{R}^{n\times n}$ by

$$U := (M - M_d) W_d, \tag{3.1}$$

where, as defined in the previous section, M is the mean first passage matrix of T and $W := e\pi^t$. As we shall see, U turns out to play a central role in establishing the connection between $A^{\#}$ and quantities related to Perron complementation and in constructing our parallel algorithm for computing $A^{\#}$. It should be noted that U is known in the literature to have several interpretations concerning Markov chains and concerning surfing the Web. First, Kemeny and Snell [17, Theorem 4.4.10] observed that U has constant row sums $\sum_{j\neq i} \pi_j m_{i,j} = \mathrm{tr}(A^{\#})$, for $i = 1, 2, \ldots, n$. The quantity $K := \mathrm{tr}(A^{\#}) + 1$ is known as the *Kemeny constant.* In relation to the Web it measures the expected number of links a surfer would follow to exit a site, which he may have

Electronic Journal of Linear Algebra ISSN 1081-3810
A publication of the International Linear Algebra Society
Volume 13, pp. 131-145, April 2005

ELA
www.math.technion.ac.il/iic/ela

entered because of being lost, and enter another random site (see Levene and Loizou [24]). It is also called the *expected time to mixing in a Markov chain* (see Hunter [16]).

The relationship between the diagonal entries of $A^{\#}$ and U is given by the lemma below, which shows that the diagonal entries of $A^{\#}$ can be expressed in terms of U as well as ξ_{α_i} and π_{α_i} arising from Perron complementation.

LEMMA 3.1. *Let* $A^{\#} = [a^{\#}_{i,j}]$. *Then*

$$[a^{\#}_{1,1}, a^{\#}_{2,2}, \ldots, a^{\#}_{n,n}] \;=\; [\xi_{\alpha_1}\pi^t_{\alpha_1}, \xi_{\alpha_2}\pi^t_{\alpha_2}, \ldots, \xi_{\alpha_k}\pi^t_{\alpha_k}]\; U. \tag{3.2}$$

Proof. First we observe that from Lemma 2.2, $m_{i,i} = 1/\pi_i$, for all i, i.e. $M_d = (W_d)^{-1}$.

Next, again by Lemma 2.2, we obtain that

$$\begin{aligned} \pi^t U &= \pi^t (M - M_d) W_d \\ &= \pi^t \Big[[I - A^{\#} + J(A^{\#})_d](W_d)^{-1} - (W_d)^{-1} \Big] W_d \\ &= -\pi^t A^{\#} + \pi^t J (A^{\#})_d \\ &= [a^{\#}_{1,1}, a^{\#}_{2,2}, \ldots, a^{\#}_{n,n}], \end{aligned}$$

where the last equality is due to Lemma 2.1. This, together with (2.5), yield (3.2). ∎

Suppose next that $A^{\#}$ and U are partitioned in conformity with T in (2.1). Our next lemma concerns the relationship between the blocks of $A^{\#}$ and those of U.

LEMMA 3.2. *For* $i = 1, 2, \ldots, k$,

$$-A^{\#}[\alpha_i] + J(A^{\#}[\alpha_i])_d = U[\alpha_i] \tag{3.3}$$

and

$$-A^{\#}[\langle n\rangle \backslash \alpha_i, \alpha_i] + J(A^{\#}[\alpha_i])_d = U[\langle n\rangle \backslash \alpha_i, \alpha_i]. \tag{3.4}$$

Proof. Similar to the proof of Lemma 3.1, we see that

$$\begin{aligned} U &= \Big[[I - A^{\#} + J(A^{\#})_d](W_d)^{-1} - (W_d)^{-1} \Big] W_d \\ &= -A^{\#} + J(A^{\#})_d. \end{aligned}$$

Without loss of generality, we consider here the case when $i = 1$. We put $\beta = \langle n\rangle \backslash \alpha_1$ for brevity. On partitioning U, $A^{\#}$, and J in conformity with T, we obtain

Electronic Journal of Linear Algebra ISSN 1081-3810
A publication of the International Linear Algebra Society
Volume 13, pp. 131-145, April 2005

ELA
www.math.technion.ac.il/iic/ela

that

$$\begin{bmatrix} U[\alpha_1] & U[\alpha_1,\beta] \\ U[\beta,\alpha_1] & U[\beta] \end{bmatrix} = -\begin{bmatrix} A^{\#}[\alpha_1] & A^{\#}[\alpha_1,\beta] \\ A^{\#}[\beta,\alpha_1] & A^{\#}[\beta] \end{bmatrix} + \begin{bmatrix} J(A^{\#}[\alpha_1])_d & J(A^{\#}[\beta])_d \\ J(A^{\#}[\alpha_1])_d & J(A^{\#}[\beta])_d \end{bmatrix},$$

from which (3.3) and (3.4) follow. □

We are now in a position to introduce our main results. For that purpose recall that P_{α_i} is the Perron complement of $T[\alpha_i]$, $B_{\alpha_i} := I - P_{\alpha_i}$, ξ_{α_i} is the i–th coupling factor, $W_{\alpha_i} := e\pi^t_{\alpha_i}$, where π_{α_i} is the normalized left Perron vector of P_{α_i}, and V_{α_i} is the skew–symmetric matrix given in (2.8). In Lemma 3.2 we showed how the blocks of $A^{\#}$ are related to the corresponding blocks of U. We shall show next how the blocks of U can be obtained from the group inverses of the B_{α_i}'s starting with the diagonal blocks of U.

THEOREM 3.3. *For $i = 1, 2, \ldots, k$,*

$$U[\alpha_i] \;=\; -B^{\#}_{\alpha_i} + J(B^{\#}_{\alpha_i})_d + \xi_{\alpha_i} V_{\alpha_i}(W_{\alpha_i})_d. \tag{3.5}$$

Proof. It suffices to show the conclusion for the case when $i = 1$. By Lemma 2.2 and on partitioning M, $A^{\#}$, and W in conformity with T, we have, using (2.5), that

$$\begin{aligned} M[\alpha_1] &= \Big[I - A^{\#}[\alpha_1] + J(A^{\#}[\alpha_1])_d\Big][(W[\alpha_1])_d]^{-1} \\ &= \frac{1}{\xi_{\alpha_1}}\Big[I - A^{\#}[\alpha_1] + J(A^{\#}[\alpha_1])_d\Big][(W_{\alpha_1})_d]^{-1}. \end{aligned} \tag{3.6}$$

On the other hand, by (2.4) and (2.7), we find that

$$M[\alpha_1] \;=\; \frac{1}{\xi_{\alpha_1}}\Big[I - B^{\#}_{\alpha_1} + J(B^{\#}_{\alpha_1})_d\Big][(W_{\alpha_1})_d]^{-1} + V_{\alpha_1}. \tag{3.7}$$

It thus follows from (3.6) and (3.7) that

$$\frac{1}{\xi_{\alpha_1}}\Big[-A^{\#}[\alpha_1] + J(A^{\#}[\alpha_1])_d\Big][(W_{\alpha_1})_d]^{-1} \;=\; \frac{1}{\xi_{\alpha_1}}\Big[-B^{\#}_{\alpha_1} + J(B^{\#}_{\alpha_1})_d\Big][(W_{\alpha_1})_d]^{-1} + V_{\alpha_1},$$

which, by (3.3), can be reduced to

$$\begin{aligned} U[\alpha_1] &= -A^{\#}[\alpha_1] + J(A^{\#}[\alpha_1])_d \\ &= -B^{\#}_{\alpha_1} + J(B^{\#}_{\alpha_1})_d + \xi_{\alpha_1} V_{\alpha_1}(W_{\alpha_1})_d. \end{aligned}$$

This completes the proof. □

Electronic Journal of Linear Algebra ISSN 1081-3810
A publication of the International Linear Algebra Society
Volume 13, pp. 131-145, April 2005

ELA
www.math.technion.ac.il/iic/ela

Having determined a representation for the diagonal blocks of U, we now determine a representation for the off–diagonal blocks of U.

THEOREM 3.4. *For $i = 1, 2, \ldots, k$,*

$$U[\langle n\rangle\backslash\alpha_i, \alpha_i] = (I - T[\langle n\rangle\backslash\alpha_i])^{-1}\Big[T[\langle n\rangle\backslash\alpha_i, \alpha_i]U[\alpha_i] + \xi_{\alpha_i} J(W_{\alpha_i})_d\Big]. \tag{3.8}$$

Proof. It suffices for the sake of simplicity to set $i = 1$ and $\beta = \langle n\rangle\backslash\alpha_1$. Thus what we need to show is the following:

$$U[\beta, \alpha_1] = (I - T[\beta])^{-1}\Big[T[\beta, \alpha_1]U[\alpha_1] + \xi_{\alpha_1} J(W_{\alpha_1})_d\Big].$$

Again, by Lemma 2.2 and on partitioning M, $A^{\#}$, and W in conformity with T, we have, using (2.5), that

$$\begin{aligned} M[\beta, \alpha_1] &= \Big[- A^{\#}[\beta, \alpha_1] + J(A^{\#}[\alpha_1])_d\Big][(W[\alpha_1])_d]^{-1} \\ &= \frac{1}{\xi_{\alpha_1}}\Big[- A^{\#}[\beta, \alpha_1] + J(A^{\#}[\alpha_1])_d\Big][(W_{\alpha_1})_d]^{-1}. \end{aligned} \tag{3.9}$$

On the other hand, by (2.4) and (2.7), we see that

$$\begin{aligned} M[\alpha_1] - (M[\alpha_1])_d &= \frac{1}{\xi_{\alpha_1}}\Big[M_{\alpha_1} - (M_{\alpha_1})_d\Big] + V_{\alpha_1} \\ &= \frac{1}{\xi_{\alpha_1}}\Big[- B^{\#}_{\alpha_1} + J(B^{\#}_{\alpha_1})_d\Big][(W_{\alpha_1})_d]^{-1} + V_{\alpha_1}, \end{aligned}$$

which, together with (2.9), yield that

$$\begin{aligned} M[\beta, \alpha_1] &= \frac{1}{\xi_{\alpha_1}}(I - T[\beta])^{-1}\Big[T[\beta, \alpha_1][-B^{\#}_{\alpha_1} + J(B^{\#}_{\alpha_1})_d][(W_{\alpha_1})_d]^{-1} \\ &\qquad + \xi_{\alpha_1} T[\beta, \alpha_1]V_{\alpha_1} + \xi_{\alpha_1} J\Big] \\ &= \frac{1}{\xi_{\alpha_1}}(I - T[\beta])^{-1}\Big[T[\beta, \alpha_1]U[\alpha_1] + \xi_{\alpha_1} J(W_{\alpha_1})_d\Big][(W_{\alpha_1})_d]^{-1}. \end{aligned} \tag{3.10}$$

The conclusion follows now from (3.4), (3.9), and (3.10). □

Theorems 3.3 and 3.4 make it clear that U can be completely determined by quantities directly related to Perron complementation. Furthermore, as shown by our next theorem, $A^{\#}$ can be completely determined by U and thus also by quantities directly related to Perron complementation.

Electronic Journal of Linear Algebra ISSN 1081-3810
A publication of the International Linear Algebra Society
Volume 13, pp. 131-145, April 2005

ELA
www.math.technion.ac.il/iic/ela

THEOREM 3.5. *For $i = 1, 2, \ldots, k$, the diagonal blocks of $A^{\#}$ are given by:*

$$\begin{aligned} A^{\#}[\alpha_i] &= e[\xi_{\alpha_1}\pi^t_{\alpha_1}, \xi_{\alpha_2}\pi^t_{\alpha_2}, \ldots, \xi_{\alpha_k}\pi^t_{\alpha_k}]\ U[\langle n\rangle, \alpha_i] - U[\alpha_i] \\ &= e\sum_{\ell=1}^{k} \xi_{\alpha_\ell}\pi^t_{\alpha_\ell}\ U[\alpha_\ell, \alpha_i] - U[\alpha_i]. \end{aligned} \tag{3.11}$$

The off–diagonal blocks of $A^{\#}$ are given as follows: for $j \neq i$, $1 \leq i, j \leq k$,

$$A^{\#}[\alpha_j, \alpha_i] \ = \ e\sum_{\ell=1}^{k} \xi_{\alpha_\ell}\pi^t_{\alpha_\ell}\ U[\alpha_\ell, \alpha_i] - U[\alpha_j, \alpha_i]. \tag{3.12}$$

Proof. From (3.3) we know that

$$A^{\#}[\alpha_i] = J(A^{\#}[\alpha_i])_d - U[\alpha_i].$$

On the other hand, from (3.2), we have that

$$J(A^{\#}[\alpha_i])_d = e[\xi_{\alpha_1}\pi^t_{\alpha_1}, \xi_{\alpha_2}\pi^t_{\alpha_2}, \ldots, \xi_{\alpha_k}\pi^t_{\alpha_k}]\ U[\langle n\rangle, \alpha_i]. \tag{3.13}$$

Thus (3.11) follows.

Similarly, using (3.4) and (3.13), we see that (3.12) holds. □

Theorems 3.3, 3.4, and 3.5 clearly illustrate how the blocks of $A^{\#}$ can be assembled from quantities arising from Perron complementation, namely from ξ_{α_i}, π_{α_i}, V_{α_i}, and $B^{\#}_{\alpha_i}$. In the next section we shall use this important fact to construct a parallel algorithm for computing $A^{\#}$.

We comment that according to [29, Formulae (7) and (19)], if α_1 is chosen to be $\{1, 2, \ldots, n-1\}$, then

$$A^{\#}[\alpha_1] \ = \ F + \xi_{\alpha_1}\Big[\mu W_{\alpha_1} - FW_{\alpha_1} - W_{\alpha_1}F\Big] \tag{3.14}$$

and

$$B^{\#}_{\alpha_1} \ = \ F + \frac{\mu}{\xi_{\alpha_1}} W_{\alpha_1} - FW_{\alpha_1} - W_{\alpha_1}F, \tag{3.15}$$

where $F = (I - T[\alpha_1])^{-1}$ and where $\mu = \xi_{\alpha_1}\pi^t_{\alpha_1}Fe$. This result, which implies the close connection between $A^{\#}[\alpha_1]$ and $B^{\#}_{\alpha_1}$, can be regarded as a special case of our Theorems 3.3, 3.4, and 3.5.

Consider the case when T is not only irreducible and stochastic, but also d–periodic, that is,

$$T = \begin{bmatrix} 0 & T_1 & 0 & \cdots & \cdots & 0 \\ 0 & 0 & T_2 & 0 & \cdots & 0 \\ \vdots & \vdots & \ddots & \ddots & \ddots & \vdots \\ \vdots & \vdots & & \ddots & \ddots & 0 \\ 0 & 0 & \cdots & \cdots & \ddots & T_{d-1} \\ T_d & 0 & \cdots & \cdots & \cdots & 0 \end{bmatrix}, \tag{3.16}$$

Electronic Journal of Linear Algebra ISSN 1081-3810
A publication of the International Linear Algebra Society
Volume 13, pp. 131-145, April 2005

where the diagonal zero blocks are square and where $(I-T)^{\#}$ is partitioned conformably with T in (3.16). For this case Kirkland [18, Theorem 1] has shown how to compute the blocks of $(I-T)^{\#}$ in terms of $(I-P_j)^{\#}$, where $P_1 := T_1T_2\cdots T_d$ and where, for $j = 2,3,\ldots,d$, $P_j := T_jT_{j+1}\cdots T_dT_1\cdots T_{j-1}$. It can be readily verified that for each j, P_j is indeed the Perron complement of the j–th diagonal block of T. This result can also be regarded as a particular case of our Theorems 3.3, 3.4, and 3.5.

We finally remark that Theorems 3.3, 3.4, and 3.5 apply to the case when any α_i consists of nonconsecutive, but increasing, indices, provided that $\alpha_1,\alpha_2,\ldots,\alpha_k$ form a partition of $\langle n\rangle$.

4. Parallel Algorithm. In this section we shall provide a parallel algorithm for the efficient computation of $A^{\#}$ via Perron complementation. The algorithm will be illustrated mainly by partitioning T into a 2×2 block matrix as follows:

$$T = \begin{bmatrix} T[\alpha_1] & T[\alpha_1,\alpha_2] \\ T[\alpha_2,\alpha_1] & T[\alpha_2] \end{bmatrix}, \tag{4.1}$$

where α_1 and α_2, both nonempty, form a partition of $\langle n\rangle$.

We begin the procedure by computing the Perron complements P_{α_1} and P_{α_2}. According to (1.2), this requires certain matrix inversions. We retain the results of these inversions since they will be needed later in the computation of V_{α_1}, V_{α_2}, and in the computation of the off–diagonal blocks of U (see (2.8) and (3.8), respectively) too. As shown in the proof of Theorem 3.5, the blocks of $A^{\#}$ are determined by the blocks of U, the coupling factors ξ_{α_1} and ξ_{α_2}, and the normalized left Perron vectors π_{α_1} and π_{α_2}. Among these quantities, each π_{α_i} can be computed separately from its respective P_{α_i}, while ξ_{α_1} and ξ_{α_2} can be obtained from the coupling matrix C once π_{α_1} and π_{α_2} are known.

Continuing, concerning the blocks of U, each diagonal block $U[\alpha_i]$ is determined by $B^{\#}_{\alpha_i}$, V_{α_i}, ξ_{α_i}, and W_{α_i} according to (3.5). Clearly each $B^{\#}_{\alpha_i}$ can be computed separately from the corresponding P_{α_i}. As soon as $B^{\#}_{\alpha_i}$ becomes available, V_{α_i} can be obtained immediately from (2.8). Having found $U[\alpha_i]$, we proceed with (3.8) so as to determine $U[\langle n\rangle\backslash\alpha_i,\alpha_i]$.

In summary, with a 2×2 partitioning scheme and Perron complementation, the computation of $A^{\#}$ can be implemented by computing the following components roughly in the order as they appear: (i) P_{α_1} and P_{α_2}; (ii) π_{α_1} and π_{α_2}; (iii) $B^{\#}_{\alpha_1}$ and $B^{\#}_{\alpha_2}$; (iv) C, ξ_{α_1}, and ξ_{α_2}; (v) V_{α_1} and V_{α_2}; (vi) $U[\alpha_1]$ and $U[\alpha_2]$; and (vii) $U[\alpha_2,\alpha_1]$ and $U[\alpha_1,\alpha_2]$. Each of these components, except (iv), consists of two separate subcomponents that can be computed concurrently and in parallel. In addition, steps (ii) and (iii) may also be executed in parallel. The flow of the computation on each P_{α_i}, as illustrated in Figure 4.1, can obviously be carried out independently, with the computation of ξ_{α_1} and ξ_{α_2} being the only exception.

According to our Theorems 3.3, 3.4, and 3.5, the above algorithm can be extended to the general $k\times k$ $(k>2)$ partitioning scheme where T is partitioned as in (2.1), which would be an appropriate choice when the computation of $A^{\#}$ is to be distributed over k processors. Specifically, we begin with the computation of the Perron

Electronic Journal of Linear Algebra ISSN 1081-3810
A publication of the International Linear Algebra Society
Volume 13, pp. 131-145, April 2005

ELA
www.math.technion.ac.il/iic/ela

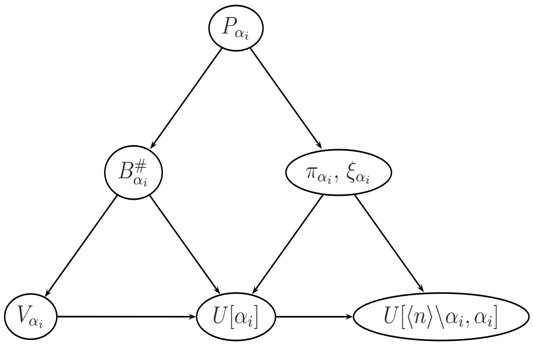

FIG. 4.1. *Computation on the Perron Complement* P_{α_i}

complements P_{α_i} for $i = 1, 2, \ldots, k$, then for each i we proceed to calculating $B^{\#}_{\alpha_i}$, π_{α_i} and ξ_{α_i}, V_{α_i}, $U[\alpha_i]$, and $U[\langle n\rangle\backslash\alpha_i, \alpha_i]$ in the order as shown in Figure 4.1, and finally we recover the blocks of $A^{\#}$ using formulae (3.11) and (3.12). Alternatively, similar to what was suggested in [27], a higher level partitioning scheme can also be achieved by following successive lower level partitionings on the Perron complements. For example, we may partition T into a 2×2 block form and construct the Perron complements P_{α_1} and P_{α_2}, but then for each i, we may compute the group inverse $B^{\#}_{\alpha_i}$ using our parallel algorithm by further partitioning the corresponding P_{α_i} into a 2×2 block form. In other words, we may compute $B^{\#}_{\alpha_i}$ on the Perron complements of P_{α_i}, rather than directly computing it on P_{α_i}.

It is natural to ask whether the above parallel algorithm for computing $A^{\#}$ via Perron complementation is less costly than the direct computation of $A^{\#}$, considering that there are quite a few quantities, though all are smaller than n in size, to compute. To answer this question, we shall estimate the asymptotic number of multiplications necessary to implement the parallel algorithm and compare it with the case without parallelism.

It should be mentioned that for general matrices, there are actually very few numerically viable methods for computing $A^{\#}$ because of the issue of numerical instability. Specifically, problems might arise in determining the bases for $R(A)$ and $N(A)$ ([2, 4]) or the characteristic polynomial of A and the resolvent $(zI - A)^{-1}$ ([11, 13]). Under our assumption on A, namely that A is an irreducible singular M–matrix, however, there is a quite reliable method for computing $A^{\#}$ proposed by Meyer [25, Section 5], which can be implemented using the conventional Gaussian eliminations on nonsingular matrices. Here we shall adopt this method while counting the asymptotic numbers of multiplications.

According to Meyer [25, Theorem 5.5 and Table 2], if $A^{\#}$ is computed directly from the formula

$$A^{\#} = (A + W)^{-1} - W \tag{4.2}$$

and if π is computed first, then the number of multiplications required is roughly $4n^3/3$.

Electronic Journal of Linear Algebra ISSN 1081-3810
A publication of the International Linear Algebra Society
Volume 13, pp. 131-145, April 2005

ELA
www.math.technion.ac.il/iic/ela

Suppose now that $n = 2m$ and that $\alpha_1 = \{1, 2, \ldots, m\}$ and $\alpha_2 = \{m+1, m+2, \ldots, n\}$. As shown in Figure 4.1, the total number of multiplications required for the computation on each P_{α_i} can be counted as follows:

(i) P_{α_i}. To find $(I-T[\langle n\rangle\backslash\alpha_i])^{-1}T[\langle n\rangle\backslash\alpha_i, \alpha_i]$, it requires $4m^3/3$ multiplications to solve the matrix equation $(I-T[\langle n\rangle\backslash\alpha_i])X = T[\langle n\rangle\backslash\alpha_i, \alpha_i]$ with Gaussian eliminations. An additional amount of m^3 multiplications is then needed to premultiply $(I-T[\langle n\rangle\backslash\alpha_i])^{-1}T[\langle n\rangle\backslash\alpha_i, \alpha_i]$ by $T[\alpha_i, \langle n\rangle\backslash\alpha_i]$.

(ii) π_{α_i} and $B^{\#}_{\alpha_i}$. Based on the result in [25], it requires $4m^3/3$ multiplications if π_{α_i} is computed first and then $B^{\#}_{\alpha_i}$ is computed from a formula similar to (4.2). We comment that the computation of ξ_{α_1} and ξ_{α_2} is trivial and does not require $O(m^3)$ multiplications.

(iii) V_{α_i}. Due to the particular structure of V_{α_i}, it is enough to compute the first column of the rank–one matrix $B^{\#}_{\alpha_i}T[\alpha_i, \langle n\rangle\backslash\alpha_i](I-T[\langle n\rangle\backslash\alpha_i])^{-1}J$. This does not require $O(m^3)$ multiplications since the LU factorization in (i) can be retained.

(iv) $U[\alpha_i]$. Clearly this does not require $O(m^3)$ multiplications.

(v) $U[\langle n\rangle\backslash\alpha_i, \alpha_i]$. The computation of $\xi_{\alpha_i}(I-T[\langle n\rangle\backslash\alpha_i])^{-1}J(W_{\alpha_i})_d$ does not require $O(m^3)$ multiplications for exactly the same reason as mentioned in (iii). It requires m^3 multiplications, however, to postmultiply $(I-T[\langle n\rangle\backslash\alpha_i])^{-1}\times T[\langle n\rangle\backslash\alpha_i, \alpha_i]$ by $U[\alpha_i]$.

(vi) $A^{\#}[\alpha_i]$ and $A^{\#}[\langle n\rangle\backslash\alpha_i, \alpha_i]$. These can be obtained without multiplication operations from (3.3) and (3.4), respectively, when $J(A^{\#}[\alpha_i])_d$ is known. Note that, according to (3.13), the computation of $J(A^{\#}[\alpha_i])_d$ does not require $O(m^3)$ multiplications.

From (i) through (vi), we conclude that the number of multiplications necessary for implementing the parallel algorithm for computing $A^{\#}$ is roughly $28m^3/3 = 7n^3/6$. On the other hand, to compute $A^{\#}$ directly from (4.2) requires $4n^3/3 = 8n^3/6$ multiplications. Therefore the parallel algorithm actually saves approximately 1/8 or 12.5% of multiplication operations.

n	flops$_d$	flops$_p$	(flops$_d$−flops$_p$)/flops$_d$
50	3,489,825	1,796,824	48.5%
100	28,126,945	14,066,448	50.0%
200	221,017,525	108,340,394	51.0%
400	1.7449×10^9	8.4711×10^8	51.5%
800	1.3551×10^{10}	6.7867×10^9	49.9%
1600	1.0575×10^{11}	5.3352×10^{10}	49.5%

Table 1: Results of Numerical Tests

We have tested with MATLAB several examples using randomly generated dense matrices and counted the number of flops with MATLAB's built–in function `flops`. The results are given in Table 1, where flops$_p$ stands for the total number of flops used in the parallel algorithm for computing $A^{\#}$, while flops$_d$ stands for that from directly computing $A^{\#}$ by (4.2).

The data in Table 1 confirms the efficiency of the parallel algorithm in terms of

Electronic Journal of Linear Algebra ISSN 1081-3810
A publication of the International Linear Algebra Society
Volume 13, pp. 131-145, April 2005
ELA
www.math.technion.ac.il/iic/ela

a reduced amount of computation[4]. The savings shown in the table are much greater than the asymptotic estimate of a 12.5% reduction in the number of multiplications. The reason appears to be that we have only tested matrices of small to moderate size. There may be also some dependency on the manner in which MATLAB performs matrix operations (multiplications and additions) at machine level.

Finally we make a few comments on the numerical accuracy of our parallel algorithm. Our parallel algorithm basically involves two main tasks: to compute the Perron complements P_{α_i} and to compute the group inverses $B_{\alpha_i}^{\#}$ associated with those Perron complements. For the same reason as we mentioned earlier, Meyer's method in [25] is advisable for computing the $B_{\alpha_i}^{\#}$'s. This method can be carried out with Gaussian eliminations as an inversion algorithm on the nonsingular matrices $B_{\alpha_i} + W_{\alpha_i}$ and therefore the standard results on round–off analysis of Gaussian eliminations (see, for example, [10]) apply. It should be noted that compared with that of $A + W$ in (4.2), the inversion of $B_{\alpha_i} + W_{\alpha_i}$ tends to be more stable because of the reduced size of the problem [10, Theorem 3.3.1]. In addition, according to [29], the computation of π_{α_i} on P_{α_i} tends to be more stable than that of π on T; in particular, the bound on the relative errors in π_{α_i} does not exceed that in π. On the other hand, Gaussian eliminations can also be used to invert the matrices $I - T[\alpha_i]$ arising from Perron complementation. Even though such nonsingular principal submatrices of $I - T$ could be poorly conditioned [8], various partitioning schemes may be exploited so as to alleviate possible numerical difficulties in calculating $(I - T[\alpha_i])^{-1}$. Consider, for example, the following 4×4 irreducible stochastic matrix:

$$T = \begin{bmatrix} .4332 & .5667 & .0001 & .0000 \\ .4331 & .5668 & .0000 & .0001 \\ .0000 & .0001 & .3667 & .6332 \\ .0001 & .0000 & .3668 & .6331 \end{bmatrix}.$$

Using the condition number $\kappa_\infty(X) := \|X\|_\infty \|X^{-1}\|_\infty$, where $X \in \mathbb{R}^{n \times n}$ is nonsingular, we obtain that for $\alpha_1 = \{1, 2\}$ and $\alpha_2 = \{3, 4\}$, $\kappa_\infty(I - T[\alpha_1]) \approx 1.1335 \times 10^4$ and $\kappa_\infty(I - T[\alpha_2]) \approx 1.2665 \times 10^4$, but for $\alpha_1 = \{1, 3\}$ and $\alpha_2 = \{2, 4\}$, $\kappa_\infty(I - T[\alpha_1]) \approx 1.1175$ and $\kappa_\infty(I - T[\alpha_2]) \approx 1.1810$.

5. Concluding Remarks. The goal of this paper was to present an efficient parallel algorithm for computing the group inverse of the singular M–matrix $A = I - T$, where T is an irreducible stochastic matrix, via Perron complementation. This algorithm can be easily modified to handle the more general case that $A = \rho_S I - S$, where S is an irreducible nonnegative matrix and where ρ_S is the Perron root of S.

As shown in Theorems 3.3, 3.4, and 3.5, the group inverse of A is closely related to the group inverses associated with the Perron complements of T. This adds to previous computational utilization of Perron complementation due to Meyer [27, 28] and Kirkland, Neumann, and Xu [22]. It remains an interesting question whether the Perron complementation approach is applicable to other computational problems relating to irreducible stochastic matrices.

[4] We remark, however, that the flop–count in the table also reflects the fact that our experiments were carried out using a particular, in this case MATLAB, programming language.

Electronic Journal of Linear Algebra ISSN 1081-3810
A publication of the International Linear Algebra Society
Volume 13, pp. 131-145, April 2005

ELA
www.math.technion.ac.il/iic/ela

In this paper we have focused on implementing the parallel algorithm for computing the group inverse of A with a 2×2 partitioning scheme. We therefore remark that even though any 2×2 block–partitioning may be used, numerically it is more efficient to choose α_i and $\langle n\rangle\backslash\alpha_i$ of roughly the same size since it balances the workload between the processors. To see this note that as the size of α_i decreases, it is less costly to compute $(I - T[\alpha_i])^{-1}$, but at the same time, the size of $\langle n\rangle\backslash\alpha_i$ increases accordingly, and therefore it is more costly to compute $(I - T[\langle n\rangle\backslash\alpha_i])^{-1}$.

The operational count presented in Table 1 shows that the parallel algorithm is capable of significantly reducing the amount of necessary multiplication operations as compared with directly computing the group inverse of A. It is interesting to observe that in [22], the computation of the mean first passage matrix of a finite ergodic Markov chain with transition matrix T is carried out in parallel on the Perron complements of T. A crucial step there is the computation of the group inverse associated with each Perron complement. When this step is accomplished with the parallel algorithm for computing $A^{\#}$ developed here, we can expect that in the parallel computation of the mean first passage matrix as suggested in [22], further reductions in the computational effort can be achieved.

REFERENCES

[1] K. M. Anstreicher and U. G. Rothblum. Using Gauss–Jordan elimination to compute the index, generalized nullspaces, and Drazin inverses. *Lin. Alg. Appl.*, 85:221–239, 1987.

[2] A. Ben–Israel and T. N. E. Greville. *Generalized Inverses: Theory and Applications, 2nd ed.* Springer, New York, 2003.

[3] A. Berman and R. J. Plemmons. *Nonnegative Matrices in the Mathematical Sciences.* SIAM, Philadelphia, 1994.

[4] S. L. Campbell and C. D. J. Meyer. *Generalized Inverses of Linear Transformations.* Dover Publications, New York, 1991.

[5] E. Deutsch and M. Neumann. Derivatives of the Perron root at an essentially nonnegative matrix and the group inverse of an M–matrix. *J. Math. Anal. Appl.*, 102:1–29, 1984.

[6] E. Deutsch and M. Neumann. On the first and second order derivatives of the Perron vector. *Lin. Alg. Appl.*, 71:57–76, 1985.

[7] M. Eiermann, I. Marek, and W. Niethammer. On the solution of singular linear equations by semiiterative equations. *Numer. Math.*, 53:265–283, 1988.

[8] R. Funderlic and C. D. Meyer. Sensitivity of the stationary distribution vector for an ergodic Markov chain. *Lin. Alg. Appl.*, 76:1–17, 1986.

[9] G. Golub and C. D. Meyer. Using the QR factorization and group inversion to compute, differentiate, and estimate the sensitivity of stationary probabilities for Markov chains. *SIAM J. Alg. Disc. Meth.*, 7:273–281, 1986.

[10] G. Golub and C. Van Loan. *Matrix Computation, 3rd ed.* The John Hopkins Univ. Press, Baltimore, Maryland, 1996.

[11] T. N. E. Greville. The Souriau–Frame algorithm and the Drazin pseudoinverse. *MRC. Tech. Sum. Rep.*, 1182, 1972.

[12] M. Hanke and M. Neumann. Preconditionings and splitting for singular systems. *Numer. Math.*, 57:85–89, 1990.

[13] R. E. Hartwig. More on the Souriau–Frame algorithm and the Drazin inverses. *SIAM J. Appl. Math.*, 31:42–46, 1976.

[14] M. Haviv and L. Van der Heyden. Perturbation bounds for the stationary probabilities of a finite Markov chain. *Adv. Appl. Prob.*, 16:804–818, 1984.

[15] R. A. Horn and C. R. Johnson. *Matrix Analysis.* Cambridge Univ. Press, New York, 1985.

Electronic Journal of Linear Algebra ISSN 1081-3810
A publication of the International Linear Algebra Society
Volume 13, pp. 131-145, April 2005

ELA
www.math.technion.ac.il/iic/ela

[16] J. J. Hunter. Mixing times with applications to perturbed Markov chains. *Preprint, Massey Univ., Auckland, New Zealand*, 2003.

[17] J. G. Kemeny and J. L. Snell. *Finite Markov Chains.* Van Nostrand, Princeton, New Jersey, 1960.

[18] S. J. Kirkland. The group inverse associated with an irreducible periodic nonnegative matrix. *SIAM J. Matrix Anal. Appl.*, 16:1127–1134, 1995.

[19] S. J. Kirkland and M. Neumann. Convexity and concavity of the Perron root and vector of Leslie matrices with applications to a population model. *SIAM J. Matrix Anal. Appl.*, 15:1092–1107, 1994.

[20] S. J. Kirkland and M. Neumann. Cutpoint decoupling and first passage times for random walks on graphs. *SIAM J. Matrix Anal. Appl.*, 20:860–870, 1999.

[21] S. J. Kirkland, M. Neumann, and B. Shader. Application of Paz's inequality to perturbation bounds for Markov chains. *Lin. Alg. Appl.*, 268:183–196, 1998.

[22] S. J. Kirkland, M. Neumann, and J. Xu. A divide and conquer approach to computing the mean first passage matrix for Markov chains via Perron complement reduction. *Numer. Lin. Alg. Appl.*, 8:287–295, 2001.

[23] S. J. Kirkland, M. Neumann, and J. Xu. Convexity and elasticity of the growth rate in size–classified population models. *SIAM J. Matrix Anal. Appl.*, 26:170–185, 2004.

[24] M. Levene and G. Loizou. Kemeny's constant and the random surfer. *Amer. Math. Month.*, 109:741–745, 2002.

[25] C. D. J. Meyer. The role of the group generalized inverse in the theory of finite Markov chains. *SIAM Rev.*, 17:443–464, 1975.

[26] C. D. Meyer. The condition of a finite Markov chain and perturbation bounds for the limiting probabilities. *SIAM J. Alg. Disc. Meth.*, 1:273–283, 1980.

[27] C. D. Meyer. Uncoupling the Perron eigenvector problem. *Lin. Alg. Appl.*, 114/115:69–94, 1989.

[28] C. D. Meyer. Stochastic complementation, uncoupling Markov chains, and the theory of nearly reducible systems. *SIAM Rev.*, 31:240–272, 1989.

[29] M. Neumann and J. Xu. On the stability of the computation of the stationary probabilities of Markov chains using Perron complements. *Numer. Lin. Alg. Appl.*, 10:603–618, 2003.

[30] E. Seneta. Sensitivity analysis, ergodicity coefficients, and rank–one updates for finite Markov chains. (W. J. Stewart, ed.) *Numerical Solution of Markov Chains*, Marcel Dekker, New York, 1991.