

Technical Report

TR-2010-024

Restricted Additive Schwarz Methods for Markov Chains

by

Michele Benzi, Verena Kuhlemann

MATHEMATICS AND COMPUTER SCIENCE

EMORY UNIVERSITY

Restricted Additive Schwarz Methods for Markov Chains

Michele Benzi* Verena Kuhlemann

Department of Mathematics and Computer Science, Emory University, Atlanta, GA 30322, U.S.A.

SUMMARY

The Restricted Additive Schwarz (RAS) method is adapted to the problem of computing the stationary probability distribution vector of large, sparse, irreducible stochastic matrices. Inexact and two-level variants are also considered, as well as acceleration by Krylov subspace methods. The convergence properties are analyzed and extensive numerical experiments aimed at assessing the effect of varying the number of subdomains and the amount of overlap are discussed. Copyright © 2000 John Wiley & Sons, Ltd.

KEY WORDS: Iterative methods, preconditioning, Schwarz methods, Markov chains, GMRES

1. INTRODUCTION

Domain decomposition methods are widely used for solving large-scale linear systems of equations arising from the discretization of partial differential equations (PDEs) on parallel computers [22]. Of particular importance among domain decomposition schemes are variants of the additive and multiplicative Schwarz methods, with or without overlap. These algorithms can be cast as stationary iterations associated with matrix splittings, making a purely algebraic analysis possible; see, e.g., [11, 4]. Although there are situations where Schwarz-type methods exhibit rapid convergence, they are usually more efficient when used as preconditioners for Krylov subspace methods; see [17]. In this paper we investigate the application of various additive Schwarz methods as solvers and as preconditioners for computing the stationary probability distribution vector of ergodic Markov chains with large and sparse transition matrices. In particular, we give the first analysis and implementation of the *restricted additive Schwarz* (RAS) method in the Markov chain context. The RAS method, first proposed in [8] as a preconditioner for nonsingular linear systems $Ax = b$, is a variant of the additive Schwarz method that requires less communication and is therefore better suited for parallel implementation. Somewhat surprisingly, this method tends to exhibit convergence rates that are no worse and often better than those of the standard additive Schwarz method; see [10] for an algebraic analysis in the nonsingular case. We note that RAS is the default Schwarz

*Correspondence to:

Email addresses: benzi@mathcs.emory.edu, vkuhlem@emory.edu

preconditioner for linear systems in the popular PETSc package [3]. Additive Schwarz methods (but not RAS) for Markov chains have been analyzed in [7, 14]. The focus of these papers is primarily theoretical, and the use of Krylov subspace acceleration is not considered. Here we present some analysis and the results of computational experiments on realistic Markov models from the MARCA collection [21]. Our results show that RAS preconditioning (with inexact subdomain solves) is a promising approach to the solution of large, sparse Markov chain models.

2. DEFINITIONS AND AUXILIARY RESULTS

First we introduce some terminology and notation. If B is a matrix, we write $B \geq 0$ ($B > 0$) if its entries are nonnegative (positive). With $\sigma(A)$ we denote the spectrum and with $R(A)$ the range of A . For $\lambda \in \sigma(A)$ the *index* of A with respect to λ is the smallest integer k for which $R((\lambda I - A)^{k+1}) = R((\lambda I - A)^k)$. The index is denoted with $\text{ind}_\lambda A$. A $n \times n$ matrix A is an *M-matrix* if $A = \beta I - B$, where I is the $n \times n$ identity, $B \geq 0$ and $\rho(B) \leq \beta$, where $\rho(B)$ is the spectral radius of B . If $\rho(B) < \beta$, then A is a nonsingular M-matrix, and $A^{-1} \geq 0$. Otherwise, A is singular. Recall that A is *irreducible* if the directed graph associated with A is strongly connected.

We are interested in computing the stationary probability distribution vector of finite, ergodic Markov chains. This amounts to finding a non-trivial solution to a homogeneous system of linear equations $Ax = 0$, where A is a singular, irreducible M-matrix. It is well-known that up to normalization, such solution vector x is unique [20]. This formulation applies to Discrete-Time (DT) as well as to Continuous-Time (CT) Markov Chains. For a DTMC, we have $A = I - P^T$ where P is the row-stochastic matrix of transition probabilities. Recall that $A = M - N$ is a *splitting* if M is nonsingular. The splitting is *regular* if $M^{-1} \geq 0$ and $N \geq 0$; it is *weak regular* if $M^{-1} \geq 0$ and $M^{-1}N \geq 0$; and it is *nonnegative* if $M^{-1} \geq 0$, $M^{-1}N \geq 0$, and $NM^{-1} \geq 0$. The standard stationary iteration associated with the splitting $A = M - N$ is of the form $x^{k+1} = Tx^k + c$, $k = 0, 1, \dots$, where $T = M^{-1}N$, c is a vector, and x^0 is arbitrary. It converges for all choices of x^0 if either $\rho(T) < 1$ (in which case $\lim_{k \rightarrow \infty} T^k = 0$), or $\rho(T) = 1$ and T is *convergent*, i.e., $\lim_{k \rightarrow \infty} T^k$ exists. In the cases of interest to us, T will have unit spectral radius. Such a matrix is convergent if and only if $\text{ind}_1 T = 1$ and $\lambda \in \rho(T) \setminus \{1\}$ implies $|\lambda| < 1$. For $T \geq 0$, the latter condition can be replaced with T having positive diagonal entries [1]. Also observe that in the Markov chain context we have $c = 0$, and that the sequence x^k converges to a nontrivial solution of $Ax = 0$ only if $x^0 \notin R(M^{-1}A)$. Finally, we will need the following lemma from [6] for our analysis.

Lemma 2.1. *Let T be a nonnegative square matrix such that $Tv \leq \alpha v$ with $v > 0$. Then $\rho(T) \leq \alpha$. If, furthermore, $\rho(T) = \alpha$, then $\text{ind}_\alpha T = 1$.*

3. ALGEBRAIC FORMULATION OF RAS

We adopt the notation and the approach developed in [7, 14]. Consider the set of indices $S = \{1, \dots, n\}$ and let $S = \cup_{i=1}^K S_{i,0}$ be a partition of S into K disjoint, non-empty subsets. For each $S_{i,0}$ consider $S_{i,\delta}$ with $S_{i,0} \subseteq S_{i,\delta}$. Thus, $S = \cup_{i=1}^K S_{i,\delta}$ for all values of δ , with not

necessarily pairwise disjoint $S_{i,\delta}$. For $\delta > 1$ this notation introduces overlap. One way to find this overlap is by considering the underlying undirected graph of A . If A does not have a symmetric pattern, we use the graph of $A + A^T$ instead. To each set of nodes $S_{i,0}$ we add all nodes with distance at most δ in the underlying graph. Here, as usual, the distance is defined as the length of the shortest path in the graph connecting a given node to any node in $S_{i,0}$. The matrix $R_{i,\delta}$ is a restriction operator from the whole space \mathbf{R}^n to the subspace defined by $S_{i,\delta}$. Let $\pi_{i,\delta}$ be the matrix representation of the permutation that relabels the states in $S_{i,\delta}$ as $1, \dots, n_{i,\delta}$, where $n_{i,\delta} = |S_{i,\delta}|$. Then $R_{i,\delta}$ is an $n_{i,\delta} \times n$ matrix given by $R_{i,\delta} = [I_{i,\delta}|0]\pi_{i,\delta}$, where $I_{i,\delta}$ is the $n_{i,\delta} \times n_{i,\delta}$ identity matrix. The matrix

$$A_{i,\delta} = R_{i,\delta} A R_{i,\delta}^T \quad (3.1)$$

is the restriction of A to the subspace corresponding to $S_{i,\delta}$. Thus, $A_{i,\delta}$ is an $n_{i,\delta} \times n_{i,\delta}$ principal submatrix of A . If A is an irreducible singular M-matrix, $A_{i,\delta}$ is a nonsingular M-matrix [20]. Define $E_{i,0}$ as

$$E_{i,0} = R_{i,0}^T R_{i,0} = \pi_{i,0}^T \begin{pmatrix} I_{i,0} & 0 \\ 0 & 0 \end{pmatrix} \pi_{i,0}. \quad (3.2)$$

$E_{i,0}$ is an $n \times n$ diagonal matrix with a one on the diagonal for every row where $R_{i,0}^T$ has a one, and zeroes otherwise. Thus, $\sum_{i=1}^K E_{i,0} = I$. The additive Schwarz method can be given in the form of a stationary iteration, $x^{(k+1)} = T x^{(k)} + c$, where

$$T = I - \sum_{i=1}^K R_{i,\delta}^T A_{i,\delta}^{-1} R_{i,\delta} A \quad (3.3)$$

and c is a certain vector. Note that for the solution of Markov chains c will be zero. For the restricted additive Schwarz iteration the prolongation operator $R_{i,\delta}^T$ will be replaced by a prolongation operator that does not consider overlap. Define the ‘‘restricted’’ operator $\tilde{R}_{i,\delta}$ as

$$\tilde{R}_{i,\delta} = R_{i,\delta} E_{i,0}. \quad (3.4)$$

With this notation, the restricted additive Schwarz method has the form of a stationary iteration, $x^{(k+1)} = T x^{(k)} + c$, where $T = I - \sum_{i=1}^K \tilde{R}_{i,\delta}^T A_{i,\delta}^{-1} R_{i,\delta} A$. We will use

$$M_{RAS,\delta}^{-1} = \sum_{i=1}^K \tilde{R}_{i,\delta}^T A_{i,\delta}^{-1} R_{i,\delta} \quad (3.5)$$

as a right preconditioner for GMRES. That is, we will use GMRES to solve $A M_{RAS,\delta}^{-1} y = 0$, with $y = M_{RAS,\delta} x$. The notation $M_{RAS,\delta}^{-1}$ only makes sense if $\sum_{i=1}^K \tilde{R}_{i,\delta}^T A_{i,\delta}^{-1} R_{i,\delta}$ is a nonsingular matrix. The next theorem shows that in our case the matrix is indeed nonsingular.

Theorem 3.1. *If A is an irreducible nonsingular M-matrix, then $M = \sum_{i=1}^K \tilde{R}_{i,\delta}^T A_{i,\delta}^{-1} R_{i,\delta}$ is nonsingular.*

Proof We want to show that $Mx \neq 0$ for $x \neq 0$. Note that

$$A_{i,\delta} = \begin{pmatrix} A_{i,0} & \hat{K}_{i,\delta} \\ \hat{L}_{i,\delta} & \hat{A}_{-i,\delta} \end{pmatrix}$$

with $A_{i,0}$ and $\hat{A}_{-i,\delta}$ nonsingular M-matrices, and $\hat{K}_{i,\delta} \leq 0$ and $\hat{L}_{i,\delta} \leq 0$. Since the diagonal entries of $\hat{A}_{-i,\delta}$ are at most 1, and the off-diagonal entries are nonpositive, we obtain

$$A_{i,\delta} \leq \begin{pmatrix} A_{i,0} & 0 \\ 0 & I_{n_{-i,\delta}} \end{pmatrix},$$

where $I_{n_{-i,\delta}}$ is the identity matrix of size $(n_{i,\delta} - n_{i,0}) \times (n_{i,\delta} - n_{i,0})$. Set

$$\hat{A}_{i,\delta} = \begin{pmatrix} A_{i,0} & 0 \\ 0 & I_{n_{-i,\delta}} \end{pmatrix}$$

and note that $\hat{A}_{i,\delta}$ is a nonsingular M-matrix. Thus, $0 \leq \hat{A}_{i,\delta}^{-1} \leq A_{i,\delta}^{-1}$. Let $x \in \mathbf{R}^{n_{i,\delta}}$ be nonnegative with at least one positive entry. Then $0 \leq \hat{A}_{i,\delta}^{-1}x \leq A_{i,\delta}^{-1}x$ and $R_{i,0}R_{i,\delta}^T\hat{A}_{i,\delta}^{-1}x$ has at least one positive entry. Therefore, $R_{i,0}R_{i,\delta}^TA_{i,\delta}^{-1}x$ has at least one positive entry. Similarly, if $x \in \mathbf{R}^{n_{i,\delta}}$ is nonpositive with at least one negative entry, then $R_{i,0}R_{i,\delta}^TA_{i,\delta}^{-1}x \leq 0$ with at least one negative entry. So, for $x \neq 0$, $R_{i,0}R_{i,\delta}^TA_{i,\delta}^{-1}x$ has at least one nonzero entry. For a nonzero vector $x \in \mathbf{R}^n$ there exists i such that $R_{i,0}x \neq 0$. Since $Mx = \sum_{i=1}^K R_{i,0}R_{i,\delta}^TA_{i,\delta}^{-1}R_{i,\delta}x$, it follows that $R_{i,0}Mx = R_{i,0}R_{i,\delta}^TA_{i,\delta}^{-1}R_{i,\delta}x$. Thus $R_{i,0}Mx \neq 0$, which implies $Mx \neq 0$ for $x \neq 0$. \square

For the purpose of convergence analysis a different algebraic representation of $M_{RAS,\delta}^{-1}$ is required. We can write

$$\pi_{i,\delta}A\pi_{i,\delta}^T = \begin{pmatrix} A_{i,\delta} & K_{i,\delta} \\ L_{i,\delta} & A_{-i,\delta} \end{pmatrix}, \quad (3.6)$$

where $A_{-i,\delta}$ is the principal submatrix of A “complementary” to $A_{i,\delta}$. Thus, $A_{-i,\delta}$ is a nonsingular M-matrix. Let $D_{-i,\delta} = \text{diag}(A_{-i,\delta})$ and note that since $A_{-i,\delta}$ is an irreducible nonsingular M-matrix, $D_{-i,\delta} \geq 0$ with positive diagonal entries. We construct a matrix $M_{i,\delta}$ corresponding to $R_{i,\delta}$ as follows:

$$M_{i,\delta} = \pi_i^T \begin{pmatrix} A_{i,\delta} & 0 \\ 0 & D_{-i,\delta} \end{pmatrix} \pi_i. \quad (3.7)$$

Since $A_{i,\delta}$ is nonsingular and $D_{-i,\delta}$ has positive entries on the diagonal, $M_{i,\delta}$ is invertible. It has been proven by Frommer and Szyld [10] that

$$\tilde{R}_{i,\delta}^T A_{i,\delta}^{-1} R_{i,\delta} = E_{i,0} M_{i,\delta}^{-1}. \quad (3.8)$$

Thus, $M_{RAS,\delta}^{-1}$ can be written as

$$M_{RAS,\delta}^{-1} = \sum_{i=1}^K E_{i,0} M_{i,\delta}^{-1}, \quad (3.9)$$

and we will use this representation in our analysis of the restricted additive Schwarz method for Markov chains.

4. PROPERTIES OF THE RAS SPLITTING

In this section we will show that the restricted additive Schwarz iterations are convergent for irreducible singular M-matrices. By the Perron-Frobenius Theorem (see for example [5]) and Lemma 2.1 we have that $\rho(T) = 1$ and $\text{ind}_1 T = 1$. As mentioned earlier the standard stationary iteration will converge for a convergent iteration matrix T with $\rho(T) = 1$. If $T \geq 0$, it is convergent if and only if $\text{ind}_1 T = 1$ and T has positive diagonal entries. In general, T may not have positive entries along the diagonal. In this case, a slight modification of the above algebraic representation is sufficient to ensure this property. Adding a small positive constant to the diagonal entries of $M_{i,\delta}$ makes the splittings $A = M_{i,\delta} - N_{i,\delta}$ regular, and the diagonal entries of $T_{i,\delta} = M_{i,\delta}^{-1}N_{i,\delta}$ positive [14]. The following two results show that with this modification, the splitting $A = M_{RAS,\delta} - N_{RAS,\delta}$ induced by $M_{RAS,\delta}$ is weak regular and that $M_{RAS,\delta}^{-1}N_{RAS,\delta}$ has positive diagonal entries.

Theorem 4.1. *Let A be an irreducible singular M-matrix, and let $M_{RAS,\delta}$ be given as in (3.9). Then the splitting $A = M_{RAS,\delta} - N_{RAS,\delta}$ is weak regular.*

Proof With $E_{i,0} \geq 0$ and $M_{i,\delta}^{-1} \geq 0$ follows that $M_{RAS,\delta}^{-1} = \sum_{i=1}^K E_{i,0}M_{i,\delta}^{-1} \geq 0$. Let $N_{i,\delta} = M_{i,\delta} - A$, then

$$\begin{aligned} M_{RAS,\delta}^{-1}N_{RAS,\delta} &= I - \sum_{i=1}^K \left(E_{i,0}M_{i,\delta}^{-1}A \right) \\ &= I - \sum_{i=1}^K E_{i,0} + \sum_{i=1}^K E_{i,0}M_{i,\delta}^{-1}N_{i,\delta} \\ &= \sum_{i=1}^K E_{i,0}M_{i,\delta}^{-1}N_{i,\delta}. \end{aligned}$$

The last equality holds since $\sum_{i=1}^K E_{i,0} = I$. Since $M_{i,\delta}^{-1}N_{i,\delta} \geq 0$ and $E_{i,0} \geq 0$, we obtain $M_{RAS,\delta}^{-1}N_{RAS,\delta} \geq 0$.

□

Next we will show that T has positive diagonal entries provided that $M_{i,\delta}^{-1}N_{i,\delta}$ has positive diagonal entries for all splittings $A = M_{i,\delta} - N_{i,\delta}$.

Proposition 4.2. *Let A be an irreducible singular M-matrix, and let $M_{RAS,\delta}$ and $M_{i,\delta}$ be defined as in (3.9) and (3.7). Set $N_{i,\delta} = M_{i,\delta} - A$ and $N_{RAS,\delta} = M_{RAS,\delta} - A$. Assume that $M_{i,\delta}^{-1}N_{i,\delta}$ has positive entries on the main diagonal. Then, $T = M_{RAS,\delta}^{-1}N_{RAS,\delta}$ has positive entries on the main diagonal.*

Proof As seen in the proof of Theorem 4.1,

$$M_{RAS,\delta}^{-1}N_{RAS,\delta} = \sum_{i=1}^K E_{i,0}M_{i,\delta}^{-1}N_{i,\delta}.$$

For a row in which $E_{i,0}$ is zero, the diagonal entry of $E_{i,0}M_{i,\delta}^{-1}N_{i,\delta}$ is zero. For a row in which $E_{i,0}$ has a one, the diagonal entry of $E_{i,0}M_{i,\delta}^{-1}N_{i,\delta}$ equals the diagonal entry of $M_{i,\delta}^{-1}N_{i,\delta}$ which is positive by assumption. Since $\sum_{i=1}^K E_{i,0} = I$, for each row exactly one of the $E_{i,0}$ has a one. Thus, the diagonal entries of $M_{RAS,\delta}^{-1}N_{RAS,\delta}$ are positive. \square

The following proposition describes the modification that is needed to ensure that $M_{i,\delta}^{-1}N_{i,\delta}$ has positive entries along its diagonal. The proposition can be found for example in [14].

Proposition 4.3. *Let $B \geq 0$, $B^T e = e$. Let $K > 1$ be a positive integer and let $\alpha_1, \dots, \alpha_K$ be any positive real numbers. Let $A = I - B = M_i - N_i$, $i = 1, \dots, K$, be defined by*

$$M_i = \pi_i^T \begin{pmatrix} \alpha_i I + A_{i,\delta} & 0 \\ 0 & \alpha_i I + D_{-i,\delta} \end{pmatrix} \pi_i,$$

and $N_i = M_i - A$ where π_i , A_i , and D_{-i} are defined as before. Then, the splittings $A = M_i - N_i$ are regular, and the diagonals of $T_i = M_i^{-1}N_i$ are positive, for $i = 1, \dots, K$.

Another way to ensure convergence is to replace T by $T_\alpha = \alpha I + (1 - \alpha)T$ for some $\alpha \in (0, 1)$. Such a matrix is guaranteed to be convergent if $T \geq 0$, $\rho(T) = 1$ and $\text{ind}_1(T) = 1$; see, eg, [5]. Note, however, that such modification is not necessary if Krylov subspace acceleration is used.

5. EXTENSION TO INEXACT SOLVES

In this section we extend the convergence results to the case of inexact local solves. Instead of solving the linear systems $A_{i,\delta}y_i = z_i$ we want to approximate the matrices $A_{i,\delta}$ by $\hat{A}_{i,\delta}$ so that the systems $\hat{A}_{i,\delta}y_i = z_i$ are easier to solve. The diagonal modifications mentioned above can also be regarded as a type of inexact solve. The following propositions shows that under certain assumptions the restricted additive Schwarz method is also convergent in the case of inexact local solve.

Proposition 5.1. *If $\hat{A}_{i,\delta}$ is a nonsingular M -matrix with $\hat{A}_{i,\delta} \geq A_{i,\delta}$ and*

$$\hat{M}_{i,\delta} = \pi_{i,\delta}^T \begin{pmatrix} \hat{A}_{i,\delta} & 0 \\ 0 & D_{-i,\delta} \end{pmatrix} \pi_{i,\delta},$$

then the splittings $A = \hat{M}_{i,\delta} - \hat{N}_{i,\delta}$ are weak regular.

Proof First, note that $\hat{M}_{i,\delta}^{-1}$ is nonnegative, since $\hat{A}_{i,\delta}$ has a nonnegative inverse. Second, we will consider $\hat{M}_{i,\delta}^{-1}\hat{N}_{i,\delta}$ and show that it is nonnegative. Recall that A can be written as

$$A = \pi_{i,\delta}^T \begin{pmatrix} A_{i,\delta} & K_{i,\delta} \\ L_{i,\delta} & A_{-i,\delta} \end{pmatrix} \pi_{i,\delta},$$

and it follows that

$$\hat{M}_i^{-1}\hat{N}_i = I - \pi_i^T \begin{pmatrix} \hat{A}_{i,\delta}^{-1}A_{i,\delta} & \hat{A}_{i,\delta}^{-1}K_{i,\delta} \\ D_{-i,\delta}^{-1}L_{i,\delta} & D_{-i,\delta}^{-1}A_{-i,\delta} \end{pmatrix} \pi_i.$$

With $\hat{A}_{i,\delta} \geq A_{i,\delta}$, and the fact that both matrices are M-matrices it follows that $\hat{A}_{i,\delta}^{-1} \leq A_{i,\delta}^{-1}$ and $I \geq \hat{A}_{i,\delta}^{-1} A_{i,\delta}$. Furthermore, $I - D_{-i,\delta}^{-1} A_{-i,\delta}$ is nonnegative, since $D_{-i,\delta} = \text{diag}(A_{-i,\delta})$. With $\hat{A}_{i,\delta}^{-1}$, $D_{-i,\delta}^{-1} \geq 0$, and $K_{i,\delta}$, $L_{i,\delta} \leq 0$ we can conclude that $\hat{M}_i^{-1} \hat{N}_i$ is nonnegative and that the splittings $A = \hat{M}_{i,\delta} - \hat{N}_{i,\delta}$ are weak regular.

□

Consider the splitting $A = \hat{M}_{RAS,\delta} - \hat{N}_{RAS,\delta}$ with $\hat{M}_{RAS,\delta} = \sum_{i=1}^K E_{i,0} \hat{M}_{i,\delta}^{-1}$. As seen in the proof of Theorem 4.1, $\hat{M}_{RAS,\delta}^{-1} \hat{N}_{RAS,\delta} = \sum_{i=1}^K E_{i,0} \hat{M}_{i,\delta}^{-1} \hat{N}_{i,\delta}$. From the proposition above it follows that $A = \hat{M}_{RAS,\delta} - \hat{N}_{RAS,\delta}$ is weak regular as well. Following the idea of adding a positive value to the diagonal of each $\hat{M}_{i,\delta}$ as seen in Chapter 4 will lead to positive diagonal entries in $\hat{M}_{RAS,\delta}^{-1} \hat{N}_{RAS,\delta}$. Thus, the restricted additive Schwarz method is convergent if inexact local solves are used.

We are particularly interested in the case where an incomplete LU factorization is used for approximating the solution of $A_{i,\delta} y_i = z_i$. The ILU factorization was first introduced by Varga [23] and studied by Meijerink and van der Vorst [15]. Meijerink and van der Vorst showed that for an M-Matrix A and every zero pattern $Q \subset P_n$ there is a unique lower triangular matrix L with unit diagonal, a unique upper triangular matrix U , and a matrix R with

$$\begin{aligned} l_{ij} &= 0 \text{ if } (i, j) \in Q, \\ u_{ij} &= 0 \text{ if } (i, j) \in Q, \\ r_{ij} &= 0 \text{ if } (i, j) \notin Q, \end{aligned}$$

such that the splitting $A = LU - R$ is regular. Note that Q can always be chosen such that $\text{diag}(R) = 0$. So, in our case $A_{i,\delta} = L_{i,\delta} U_{i,\delta} - R_{i,\delta}$ and we use $\hat{A}_{i,\delta} = L_{i,\delta} U_{i,\delta}$ instead. Since the splitting $A_{i,\delta} = L_{i,\delta} U_{i,\delta} - R_{i,\delta}$ is regular, $R_{i,\delta}$ is nonnegative and $\hat{A}_{i,\delta} \geq A_{i,\delta}$. Varga showed in [24] that if A is a nonsingular M-matrix and B is a matrix satisfying

$$\begin{aligned} a_{ij} &\leq b_{ij} \leq 0, & \text{for } i \neq j \text{ and} \\ 0 &< a_{ii} \leq b_{ii}, \end{aligned}$$

then B is also a nonsingular M-matrix. If the drop tolerance in the ILU factorization is sufficiently small, $\hat{A}_{i,\delta}$ will satisfy both of the above conditions. In this case $\hat{A}_{i,\delta}$ is a nonsingular M-matrix and the previous discussion shows that replacing $A_{i,\delta}$ with $\hat{A}_{i,\delta}$ will preserve the convergence results. In practice, this is seldom needed. It is never needed if Krylov acceleration is used.

6. TWO-LEVEL METHOD

The rate of convergence of the additive Schwarz method may be improved with the help of a ‘coarse grid’ correction. This correction can be formed algebraically, without any reference to an actual grid; for simplicity, however, we will use the same terminology in use in the field




Figure 1. Nonzero patterns of matrices used to construct the coarse grid correction. Left: sparsity pattern of the matrix permuted by Luby's maximal independent set reordering. Right: sparsity pattern of the corresponding Schur complement, also permuted by Luby's maximal independent set reordering. The original matrix comes from the *ncd* family.

of numerical PDEs. The correction can be applied in two ways, additively or multiplicatively; see for example [4, 9, 11, 16, 19].

In this section we describe the use of a coarse grid correction for the restricted additive Schwarz method in the case of irreducible Markov chains. The correction can either be applied additively, that is in the form

$$T = I - \theta(PA_0^{-1}R + M_{RAS}^{-1})A,$$

where $\theta > 0$ is a parameter, or multiplicatively, that is

$$T = (I - PA_0^{-1}RA)(I - M_{RAS}^{-1}A),$$

where $A_0 = RAP$. Here we consider the additive approach only. For use as a preconditioner, the parameter θ can be disregarded.

Some results were obtained by using independent sets and the Schur complement. First a maximal independent set \mathcal{F} in the underlying graph of A is found. We can partition A such that

$$\pi A \pi^T = \begin{pmatrix} A_{\mathcal{F}\mathcal{F}} & A_{\mathcal{F}\mathcal{C}} \\ A_{\mathcal{C}\mathcal{F}} & A_{\mathcal{C}\mathcal{C}} \end{pmatrix}.$$

Since \mathcal{F} is an independent set, $A_{\mathcal{F}\mathcal{F}}$ is a diagonal matrix. Note that $A_{\mathcal{F}\mathcal{F}}$ is nonsingular, since A has positive entries along its diagonal. We define the restriction operator

$$R = [-A_{\mathcal{C}\mathcal{F}}A_{\mathcal{F}\mathcal{F}}^{-1} \quad I]$$

and the prolongation operator as

$$P = \begin{bmatrix} -A_{\mathcal{F}\mathcal{F}}^{-1}A_{\mathcal{F}\mathcal{C}} \\ I \end{bmatrix}.$$

With these operators $A_0 = R(\pi A \pi^T)P$ is given by the Schur complement

$$A_0 = A_{\mathcal{C}\mathcal{C}} - A_{\mathcal{C}\mathcal{F}}A_{\mathcal{F}\mathcal{F}}^{-1}A_{\mathcal{F}\mathcal{C}}.$$

Table 6.1. Properties of the generator matrices

Matrix	Number of rows/cols	Number of nonzeros
ncd07	62,196	420,036
ncd09	129,766	883,246
mutex09	65,535	1,114,079
mutex12	263,950	4,031,310
tcomm16	13,671	67,381
tcomm20	17,081	84,211
twod06	66,049	263,169
twod10	263,169	1,050,625
reliab ₁ (m)	m^2	$5m^2 - 4m$
reliab ₂ (m)		

Note that A_0 is an M-matrix [2], and thus an ILU factorization can be used to inexactly solve a system of the form $A_0 z = y$. To reduce the size of the coarse grid correction we apply the above scheme twice. That is, a maximal independent set in the underlying graph of A_0 is found and the Schur complement of A_0 is used. A few numerical results using additive correction with above independent set method are given in the section 8 in Table 8.6 and Table 8.7. A greedy algorithm was used to find the maximal independent sets [13].

7. DESCRIPTION OF THE TEST PROBLEMS

For our numerical experiments we used the generator matrices of some Markov chain models provided in the MARCA (MARKov Chain Analyzer) collection [21]. These matrices are infinitesimal generators of CTMCs. For our purposes we converted them to the form $A = I - P^T$, with P row-stochastic. A corresponds to a discrete-time Markov chain, known as the *embedded Markov chain*; see [20, Chapter 1.4.3]. In Table 6.1 the dimensions and number of non-zeroes of our selected test matrices are shown. Each matrix is named by its family and its index in the family. The matrices from the *mutex* and *ncd* families are structurally symmetric, the matrices from the *twod* family are structurally nonsymmetric, and the matrices from the *tcomm* family are very close to being structurally symmetric. We also run some test with matrices that arise from a reliability problem. We consider a simple reliability model with two different classes of machines. We assume that each class has the same number of machines. Each machine is subject to breakdown and a subsequent repair. A state is completely specified by the ordered pair (n_1, n_2) , where n_1 denotes the number of intact machines of the first class and n_2 denotes the number of intact machines of the second class. Thus, if there are m machines in each class, the total number of possible states is $(m + 1)^2$. We order these states such that state (i, j) has index $(m + 1)(m - i) + m - j + 1$. The times between successive breakdowns and successive repairs are both exponentially distributed. The breakdown rates of class 1 machines and class 2 machines are respectively λ_1 and λ_2 . Similarly, the repair rates of the two classes of machines are μ_1 and μ_2 . The transition rate matrix for the described reliability model is then given by Q below. Here $m = 3$ machines per class where used. The diagonal elements indicated by asterisks are the negated sums of the off-diagonal elements in their corresponding rows.

$$Q = \begin{pmatrix} & 1 & 2 & 3 & 4 & 5 & 6 & 7 & 8 & 9 & 10 & 11 & 12 & 13 & 14 & 15 & 16 \\ 1 & * & 3\lambda_2 & & & 3\lambda_1 & & & & & & & & & & & & \\ 2 & \mu_2 & * & 2\lambda_2 & & & 3\lambda_1 & & & & & & & & & & & \\ 3 & & 2\mu_2 & * & \lambda_2 & & & 3\lambda_1 & & & & & & & & & & \\ 4 & & & 3\mu_2 & * & & & & 3\lambda_1 & & & & & & & & & \\ 5 & \mu_1 & & & & * & 3\lambda_2 & & & 2\lambda_1 & & & & & & & & \\ 6 & & \mu_1 & & & \mu_2 & * & 2\lambda_2 & & & 2\lambda_1 & & & & & & & \\ 7 & & & \mu_1 & & & 2\mu_2 & * & \lambda_2 & & & 2\lambda_1 & & & & & & \\ 8 & & & & \mu_1 & & & 3\mu_2 & * & & & & 2\lambda_1 & & & & & \\ 9 & & & & & 2\mu_1 & & & & * & 3\lambda_2 & & & \lambda_1 & & & & \\ 10 & & & & & & 2\mu_1 & & & \mu_2 & * & 2\lambda_2 & & & \lambda_1 & & & \\ 11 & & & & & & & 2\mu_1 & & & 2\mu_2 & * & \lambda_2 & & & \lambda_1 & & \\ 12 & & & & & & & & 2\mu_1 & & & 3\mu_2 & * & & & & \lambda_1 & \\ 13 & & & & & & & & & 3\mu_1 & & & & * & 3\lambda_2 & & & \\ 14 & & & & & & & & & & 3\mu_1 & & & \mu_2 & * & 2\lambda_2 & & \\ 15 & & & & & & & & & & & 3\mu_1 & & & 2\mu_2 & * & \lambda_2 & \\ 16 & & & & & & & & & & & & 3\mu_1 & & & 3\mu_2 & * & \end{pmatrix}.$$

Our goal is to solve the singular system $\pi Q = 0$ subject to $\|\pi\|_1 = 1$ (the normalization condition). However, one can solve the equivalent system $-Q^T \pi^T = 0$ instead. Here, the coefficient matrix is a singular irreducible M-matrix and the theory developed in the previous sections applies. From this point on, we let $A = -Q^T$ and $x = \pi^T$, so that we can use the notation introduced in the earlier sections. We tested two different reliability matrices. The first one corresponds to a reliability problem with parameters $\lambda_1 = 1$, $\lambda_2 = 0.2$, $\mu_1 = 2.5$, and $\mu_2 = 6$, while the second one corresponds to a reliability problem with parameters $\lambda_1 = 2$, $\lambda_2 = 0.9$, $\mu_1 = 0.5$, and $\mu_2 = 6$.

8. NUMERICAL EXPERIMENTS

In this section we provide the results of our numerical experiments. The implementation was done in Matlab 7.8.0 on a 2.13 GHz Intel Core 2 Duo Processor with 2GB main memory. We performed a large number of tests on numerous matrices. Here we present a selection of these results to show our overall findings. The Krylov method used was GMRES [18]. For the partitioning of our matrices we used the Metis library [12]. In our tests we used $K = 2, 8, 16, 32$, and 64 domains. Metis requires the matrices to be structurally symmetric. Therefore we applied METIS to the underlying graph of $A + A^T$ for the *twod* matrices. The amount of overlap was chosen according to the distance to the domain. For a small choice of overlap, we chose all vertices with a distance of one in the underlying graph, that is, all vertices that are connected to a vertex in the domain. The *mutex* matrices have a large separator set and choosing the overlap in the described way leads to domains that are close to the size of the matrix. For these matrices we restricted the total size of a domain, that is the partition and the overlap, by $\frac{3}{4} \cdot \frac{n}{K}$, where n is the size of the matrix and K is the number of domains. For a large choice of overlap we chose all vertices that lie within a distance of ten in the underlying graph. We factored the blocks using an incomplete LU factorization with a symmetric reverse Cuthill–McKee reordering and threshold parameter $\tau = 10^{-3}$ for the *mutex* matrices and $\tau = 10^{-4}$ for




Figure 2. Left: nonzero pattern of an ncd matrix. Right: nonzero pattern of the same matrix when four subdomains are used and the matrix is reordered accordingly by METIS.

the other matrices. Compared to exact LU factorization the timings and storage requirements were dramatically reduced, while the number of iterations remained essentially unchanged (compare Table 8.1 with 8.2 and 8.3). We found that RAS always performs as well or better than the standard additive Schwarz preconditioner in terms of number of iterations, whereas in terms of total runtime the standard additive Schwarz preconditioner tends to perform slightly better. However, AS requires more interprocessor communication and thus we expect it to be less efficient in a parallel implementation. A selection of our results is presented in the following tables. Results in bold indicate the best total runtime, that is the time to construct the preconditioner plus the time taken by the Krylov method. Runtimes and iteration counts for the MARCA examples using AS and RAS with small and large overlap are given in Table 8.2 and 8.3 respectively. A few results for the reliability problems are shown in Table 8.4 and 8.5. In Table 8.6 and 8.7 results using a 2-level RAS preconditioner are given.

Our results indicate many favorable properties of the RAS preconditioner for the numerical solution of Markov chains. In particular, we focused on the rate of convergence depending on the size of the problem, the number of domains used, and the amount of overlap chosen. First of all, the time required to construct the preconditioner, when incomplete LU factorizations are used, is very small. Also note that with an increased number of domains the size of the submatrices that require an incomplete LU factorization decreases, and thus the construction time for the preconditioner decreases. This can be best observed in the largest example of the first reliability problem ($reliab_1(1200)$, see Table 8.4). Here the time to construct the preconditioner can be cut from 155 seconds for two domains to 27 seconds for 64 domains. Furthermore, the construction expenses mainly consist of the work needed to compute the incomplete LU factorizations as the time spent to find the non-overlapping domains with the METIS software is almost negligible. The task of computing several incomplete LU factorizations of different submatrices is trivial to parallelize, and we expect an even better construction time when used in a parallel setting. Another observation that can be made from our results is that there is only a slight increase in the number of iterations when the size of the problem is increased. This property is particularly desirable for preconditioners. In some

cases the number of iterations is even decreased as the size of the problem grows larger (see for example Table 8.4). The choice of number of domains also has an influence on the convergence. As the number of domains increases so does the number of iterations. This behavior is not surprising, since the propagation of information across the computational domain takes longer for an increased number of subdomains. Note, however, that for several matrices the increase in the number of iterations is moderate. As mentioned earlier, the construction time is faster if more domains are used. In most cases the time saved during the construction is smaller than the increase in the time required by the Krylov method, and thus the overall time needed increases with the number of domains. This behavior seems to be reversed for some very large problems. Here the overall time needed is reduced for a large number of domains, since the decreased construction time outweighs the larger time needed to solve the linear system (see larger problems in both reliability models, Table 8.4 and Table 8.5). Another aspect that is worth mentioning is the effect of the amount of overlap on the convergence. While larger overlap leads to an increase in construction time, the number of iterations is decreased. The most notable reduction can be seen in the second reliability model. Here, in one case the number of iterations could be cut to 28 iterations with an increase of overlap. For a small choice of overlap 92 iterations were required. In many cases the reduction in the number of iterations induced a reduction in the overall solution time.

Finally, we comment on the results obtained with the two-level method. The number of iterations could be reduced (except for the *tcomm* examples) with our two-level approach. For example, for problems coming from the *ncd* family the number of iterations could be reduced from about 30 to about 20. It can also be seen that the properties mentioned before are more pronounced in the two-level method. That is, there is an even slower increase in number of iterations as the problem size increases. Also, the increase in number of iterations with increase number of domains is reduced. Unfortunately, the increased construction time in the two-level method almost outweighs the time saved during the iterative solve. In most cases we could only observe a very small reduction in the overall time needed, if any. In our case, a better choice of overlap seems to be more effective than a two-level method. We should mention that a number of attempts were made to find more cost-effective coarsening strategies, but we were unable to obtain better results.

Table 8.1. Results for GMRES with RAS preconditioner. K is the number of domains, ‘constr.’ the time (in seconds) needed to construct the preconditioner, ‘it’ the number of GMRES(50) iterations needed to reduce the 2-norm of the residual below 10^{-12} , ‘solve’ the time (in seconds) for the GMRES(50) iterations. For each matrix, the best overall timings are in boldface. For local solves the complete LU factorization was used.

Matrix	K	RAS preconditioned GMRES with					
		small overlap			large overlap		
		constr.	solve	it	constr.	solve	it
ncd07	2	99.10	17.54	28	221.35	52.36	27
	4	38.66	9.22	28	79.96	12.12	28
	8	21.96	8.43	28	30.73	9.67	28
	16	13.26	8.22	29	16.79	9.15	28
	32	7.95	8.50	30	7.67	7.58	29
	64	4.43	7.77	33	4.49	7.23	32
tcomm16	2	0.27	0.15	5	0.29	0.15	5
	4	0.28	0.29	5	0.33	0.16	5
	8	0.28	0.12	5	0.35	0.15	5
	16	0.24	0.11	5	0.31	0.13	5
	32	0.20	0.16	7	0.24	0.12	5
	64	0.17	0.19	8	0.19	0.14	6
twod06	2	16.94	2.95	12	20.78	3.27	12
	4	6.80	2.94	14	9.72	2.95	12
	8	4.44	2.83	15	6.42	3.12	14
	16	2.91	3.62	18	4.29	3.55	17
	32	1.94	4.23	21	2.87	3.64	18
	64	1.62	5.67	27	2.11	5.02	25
reliab ₁ (400)	2	71.75	21.52	13	86.27	23.72	13
	4	35.89	12.86	13	45.12	14.32	13
	8	14.82	17.39	21	26.32	14.28	14
	16	9.23	23.38	27	14.22	11.94	14
	32	6.17	20.92	29	9.14	9.25	15
	64	4.29	19.81	30	6.19	8.68	16
reliab ₂ (400)	2	50.51	21.95	18	59.22	21.02	16
	4	30.20	18.77	18	36.92	17.87	16
	8	12.99	29.33	31	20.76	16.60	17
	16	8.48	27.62	33	12.88	14.19	18
	32	5.65	33.06	39	8.45	13.69	20
	64	4.18	45.97	52	6.05	18.13	27

Table 8.2. Here a small amount of overlap was used. K is the number of domains, ‘constr.’ the time (in seconds) needed to construct the preconditioner, ‘it’ the number of GMRES(50) iterations needed to reduce the 2-norm of the residual below 10^{-12} , ‘solve’ the time (in seconds) for the GMRES(50) iterations. The number of GMRES(50) iterations without preconditioner is also given. For local solves ILUT was used.

Matrix	GMRES	K	GMRES preconditioned with					
			AS			RAS		
			constr.	solve	it	constr.	solve	it
ncd07	>250	2	0.52	4.96	29	0.52	5.16	29
		8	0.87	4.95	29	0.87	5.32	29
		16	0.74	5.65	30	0.74	5.68	29
		32	0.77	5.68	30	0.77	5.98	30
		64	0.83	6.67	34	0.83	6.74	33
ncd09	>250	2	1.39	11.42	28	1.39	11.71	28
		8	1.86	12.75	30	1.86	13.28	30
		16	1.51	13.81	31	1.51	13.64	30
		32	1.73	14.38	31	1.73	15.38	31
		64	1.55	15.92	34	1.55	17.20	34
mutex09	21	2	5.18	0.63	6	5.18	0.66	6
		8	2.74	1.18	11	2.74	1.02	9
		16	2.71	1.44	12	2.71	1.40	11
		32	2.39	1.63	13	2.39	1.44	11
		64	2.61	1.47	14	2.61	1.36	12
mutex12	22	2	89.15	4.09	8	89.15	4.12	7
		8	32.91	7.54	12	32.91	6.46	11
		16	20.06	7.30	13	20.06	7.02	12
		32	14.16	9.15	14	14.16	8.04	13
		64	11.55	9.85	15	11.55	9.84	14
tcomm16	93	2	0.13	0.11	5	0.13	0.11	5
		8	0.14	0.09	5	0.14	0.09	5
		16	0.12	0.08	5	0.12	0.10	5
		32	0.12	0.15	7	0.12	0.16	7
		64	0.13	0.13	7	0.13	0.16	7
tcomm20	93	2	0.17	0.13	5	0.17	0.15	5
		8	0.19	0.13	5	0.19	0.15	5
		16	0.15	0.11	5	0.15	0.12	5
		32	0.15	0.12	5	0.15	0.13	5
		64	0.16	0.18	7	0.16	0.20	7
twod06	>250	2	1.28	1.62	12	1.28	1.68	12
		8	1.15	1.93	16	1.15	2.12	16
		16	0.94	2.69	19	0.94	2.93	19
		32	0.83	3.11	21	0.83	3.32	21
		64	0.79	4.52	27	0.79	4.59	27
twod10	>250	2	6.45	8.29	16	6.45	9.17	16
		8	10.47	14.95	19	10.47	11.69	18
		16	8.07	12.28	20	8.07	12.85	19
		32	4.84	20.76	26	4.84	20.23	25
		64	4.97	27.43	31	4.97	30.69	31

Table 8.3. Here a large amount of overlap was used. K is the number of domains, ‘constr.’ the time (in seconds) needed to construct the preconditioner, ‘it’ the number of GMRES(50) iterations needed to reduce the 2-norm of the residual below 10^{-12} , ‘solve’ the time (in seconds) for the GMRES(50) iterations. For local solves ILUT was used.

Matrix	K	GMRES preconditioned with					
		AS			RAS		
		constr.	solve	it	constr.	solve	it
ncd07	2	0.79	5.10	28	0.79	5.34	28
	8	0.95	5.14	29	0.95	5.05	28
	16	0.81	5.26	29	0.81	5.71	29
	32	0.81	6.01	30	0.81	5.63	29
	64	0.81	6.45	34	0.81	6.34	32
ncd09	2	1.75	11.49	28	1.75	12.05	28
	8	2.12	12.17	29	2.12	12.90	29
	16	1.72	13.18	30	1.72	13.79	30
	32	1.77	14.03	31	1.77	15.18	31
	64	1.79	15.56	33	1.79	16.02	33
tcomm16	2	0.13	0.17	5	0.13	0.20	5
	8	0.17	0.10	5	0.17	0.11	5
	16	0.13	0.09	5	0.13	0.10	5
	32	0.14	0.10	5	0.14	0.11	5
	64	0.14	0.12	6	0.14	0.14	6
tcomm20	2	0.17	0.13	5	0.17	0.14	5
	8	0.22	0.13	5	0.22	0.15	5
	16	0.19	0.13	5	0.19	0.13	5
	32	0.18	0.13	5	0.18	0.14	5
	64	0.17	0.15	6	0.17	0.16	6
twod06	2	1.57	1.53	12	1.57	1.56	12
	8	1.46	1.87	15	1.46	1.85	14
	16	1.22	2.00	17	1.22	2.49	17
	32	1.05	3.10	19	1.05	2.94	19
	64	0.96	4.30	26	0.96	4.14	25
twod10	2	6.17	9.41	16	6.17	9.19	16
	8	13.46	12.16	18	13.46	12.00	18
	16	9.77	12.66	19	9.77	13.45	19
	32	6.33	17.04	23	6.33	18.25	23
	64	6.38	26.39	28	6.38	25.00	27

Table 8.4. The matrix *reliab1* arises from a reliability problem with parameters $\lambda_1 = 1$, $\lambda_2 = 0.2$, $\mu_1 = 2.5$, and $\mu_2 = 6$. Here K is the number of domains, ‘constr.’ the time (in seconds) needed to construct the preconditioner, ‘it’ the number of GMRES(50) iterations needed to reduce the 2-norm of the residual below 10^{-12} , ‘solve’ the time (in seconds) for the GMRES(50) iterations.

Matrix	K	RAS Preconditioned GMRES with					
		small overlap			large overlap		
		constr.	solve	it	constr.	solve	it
reliab ₁ (100)	2	0.14	0.37	15	0.16	0.23	10
	4	0.14	0.34	15	0.17	0.24	10
	8	0.15	0.31	15	0.17	0.24	10
	16	0.11	0.49	20	0.14	0.35	14
	32	0.11	0.72	24	0.13	0.60	21
	64	0.19	0.81	27	0.18	0.72	24
reliab ₁ (400)	2	4.51	4.47	13	4.76	4.50	13
	4	5.23	4.85	13	5.74	5.14	13
	8	3.62	10.38	22	4.63	5.55	14
	16	3.17	15.04	27	4.08	6.52	15
	32	2.83	17.95	29	3.51	7.17	16
	64	2.52	17.86	30	3.04	7.26	17
reliab ₁ (700)	2	21.69	53.00	26	21.75	17.52	16
	4	18.98	41.59	26	20.05	18.16	16
	8	13.33	54.13	30	15.47	19.29	16
	16	12.65	54.92	31	14.76	21.88	17
	32	9.45	50.95	30	12.06	22.88	16
	64	7.39	41.42	27	9.59	19.64	16
reliab ₁ (1000)	2	112.94	53.47	17	116.71	42.47	17
	4	75.26	42.75	17	65.91	55.57	17
	8	45.66	47.74	19	48.00	45.74	18
	16	27.23	80.44	27	29.62	42.28	18
	32	21.46	196.66	36	26.32	48.84	19
	64	19.31	251.73	36	23.82	60.82	19
reliab ₁ (1200)	2	154.85	161.76	19	160.28	157.36	19
	4	112.29	61.02	18	108.95	71.54	18
	8	73.51	224.24	33	77.90	62.63	19
	16	50.25	227.65	32	56.04	113.21	20
	32	38.74	228.14	37	43.84	73.57	20
	64	26.60	143.57	30	32.89	94.01	20

Table 8.5. The matrix *reliab2* arises from a reliability problem with parameters $\lambda_1 = 2$, $\lambda_2 = 0.9$, $\mu_1 = 0.5$, and $\mu_2 = 6$. Here K is the number of domains, ‘constr.’ the time (in seconds) needed to construct the preconditioner, ‘it’ the number of GMRES(50) iterations needed to reduce the 2-norm of the residual below 10^{-12} , ‘solve’ the time (in seconds) for the GMRES(50) iterations.

Matrix	K	RAS Preconditioned GMRES with					
		small overlap			large overlap		
		constr.	solve	it	constr.	solve	it
reliab ₂ (100)	2	0.20	0.37	14	0.24	0.34	12
	4	0.18	0.32	14	0.23	0.33	12
	8	0.17	0.47	20	0.20	0.40	14
	16	0.17	1.20	30	0.17	0.80	22
	32	0.16	1.28	35	0.20	1.04	28
	64	0.14	1.78	42	0.15	1.71	40
reliab ₂ (400)	2	7.54	10.24	19	8.27	6.83	17
	4	5.81	7.93	19	6.65	7.27	17
	8	4.21	18.85	31	5.33	8.09	18
	16	3.83	21.21	33	4.82	7.89	18
	32	2.96	31.28	39	4.11	10.19	20
	64	2.96	43.42	52	3.54	14.49	27
reliab ₂ (700)	2	50.86	50.88	25	48.79	25.99	20
	4	22.25	35.34	25	23.65	24.83	20
	8	12.41	69.96	38	13.75	28.50	22
	16	13.76	72.80	37	15.81	29.15	22
	32	10.08	121.10	49	13.62	31.67	23
	64	9.05	223.73	92	12.19	46.52	28
reliab ₂ (1000)	2	102.53	111.52	25	102.22	68.41	23
	4	80.11	72.52	25	82.25	64.44	23
	8	45.23	77.08	26	48.13	67.96	24
	16	36.72	425.31	88	41.34	76.41	26
	32	25.78	187.05	44	30.74	76.68	26
	64	20.82	272.93	69	28.14	106.39	28
reliab ₂ (1200)	2	90.33	126.93	26	92.78	98.79	25
	4	123.10	206.22	26	125.40	130.95	25
	8	81.95	146.24	27	87.95	128.44	25
	16	58.11	238.04	41	62.35	133.43	27
	32	44.42	245.89	43	49.78	124.00	27
	64	34.60	454.97	81	42.56	148.94	30

Table 8.6. A few results with a 2-level methods are shown. Here K is the number of domains, ‘constr.’ the time (in seconds) needed to construct the preconditioner, ‘it’ the number of GMRES(50) iterations needed to reduce the 2-norm of the residual below 10^{-12} , ‘solve’ the time (in seconds) for the GMRES(50) iterations.

Matrix	K	GMRES preconditioned with 2-level RAS					
		small overlap			large overlap		
		constr.	solve	it	constr.	solve	it
ncd07	2	2.31	3.01	19	2.47	3.02	19
	4	2.45	3.13	19	2.65	3.36	19
	8	2.70	3.27	19	2.80	3.24	19
	16	2.57	3.95	19	2.65	4.03	20
	32	2.57	3.84	20	2.68	3.76	20
	64	2.69	3.68	21	2.57	3.24	20
ncd09	2	5.22	7.27	19	5.57	7.09	19
	4	5.48	7.32	19	5.61	7.16	19
	8	5.55	7.46	20	5.88	7.93	20
	16	5.36	8.43	21	5.59	8.85	21
	32	5.63	9.28	21	5.69	9.30	21
	64	5.89	9.82	22	6.44	9.56	22
tcomm16	2	0.43	0.25	7	0.39	0.20	7
	4	0.39	0.18	7	0.41	0.21	7
	8	0.41	0.17	7	0.45	0.23	7
	16	0.37	0.17	7	0.45	0.19	7
	32	0.46	0.21	7	0.45	0.22	7
	64	0.43	0.24	8	0.47	0.24	7
tcomm20	2	0.58	0.34	7	0.57	0.30	7
	4	0.60	0.33	7	0.61	0.30	7
	8	0.59	0.30	7	0.64	0.31	7
	16	0.55	0.27	7	0.61	0.28	7
	32	0.54	0.28	7	0.58	0.28	7
	64	0.56	0.35	8	0.58	0.30	7
twod06	2	6.35	1.88	10	5.97	1.98	10
	4	3.19	1.79	11	3.47	1.69	10
	8	3.09	2.44	12	3.56	1.95	11
	16	2.85	2.66	14	3.19	2.77	14
	32	2.78	3.21	16	3.08	2.94	15
	64	2.71	4.46	21	3.08	4.29	20
twod10	2	22.10	8.47	12	16.65	8.21	12
	4	20.34	8.77	12	21.45	8.62	12
	8	21.54	11.02	14	24.72	10.69	13
	16	18.31	11.40	14	20.17	11.75	14
	32	15.37	17.09	18	17.31	15.84	17
	64	15.58	24.79	22	17.50	22.51	20

Table 8.7. A few results with a 2-level methods are shown. Here K is the number of domains, ‘constr.’ the time (in seconds) needed to construct the preconditioner, ‘it’ the number of GMRES(50) iterations needed to reduce the 2-norm of the residual below 10^{-12} , ‘solve’ the time (in seconds) for the GMRES(50) iterations.

Matrix	K	RAS Preconditioned GMRES with					
		small overlap			large overlap		
		constr.	solve	it	constr.	solve	it
reliab ₁ (100)	2	0.37	0.40	14	0.35	0.31	13
	4	0.34	0.34	14	0.36	0.32	13
	8	0.34	0.34	14	0.37	0.35	13
	16	0.33	0.40	16	0.34	0.38	14
	32	0.42	0.53	18	0.36	0.53	17
	64	0.41	0.58	19	0.35	0.52	18
reliab ₁ (400)	2	8.23	6.64	15	8.36	6.68	15
	4	8.55	6.88	15	8.78	6.95	15
	8	8.05	9.00	18	8.63	7.09	15
	16	7.75	11.00	20	8.31	8.04	16
	32	7.94	12.24	21	8.20	8.38	16
	64	7.58	11.85	21	8.05	8.87	17
reliab ₁ (700)	2	26.88	28.69	20	27.32	22.05	17
	4	24.18	28.64	20	24.65	22.44	17
	8	23.80	31.04	21	24.69	22.68	17
	16	23.02	33.56	22	24.12	22.97	17
	32	22.15	31.12	21	23.63	23.25	17
	64	21.70	29.53	20	23.15	23.79	17
reliab ₂ (100)	2	0.41	0.45	14	0.44	0.41	14
	4	0.38	0.39	14	0.43	0.42	14
	8	0.43	0.42	16	0.40	0.44	15
	16	0.40	0.65	21	0.44	0.61	19
	32	0.40	0.73	22	0.39	0.72	21
	64	0.42	1.00	26	0.47	1.16	25
reliab ₂ (400)	2	10.40	9.28	19	10.60	9.21	19
	4	10.26	9.13	19	10.46	9.26	19
	8	9.70	11.19	22	10.15	10.08	20
	16	9.25	12.41	23	10.10	9.77	19
	32	9.28	15.03	25	9.91	12.63	21
	64	9.67	19.29	29	9.91	15.25	24
reliab ₂ (700)	2	49.69	42.19	24	48.96	38.35	23
	4	44.34	40.44	24	45.58	37.59	23
	8	41.97	67.55	29	43.40	43.06	25
	16	40.99	48.30	27	42.00	41.73	24
	32	39.96	57.47	29	41.61	45.48	25
	64	39.81	66.60	32	41.42	51.54	27

9. SUMMARY AND CONCLUSIONS

In this paper we have extended the RAS method to the computation of the stationary vector of large, sparse Markov chains. Our results suggest that when combined with GMRES acceleration and inexact solves, RAS is a promising approach for the solution of Markov chains with large, sparse transition matrices. Future work will focus on parallel implementation and multi-level extensions.

REFERENCES

1. G. Alefeld and N. Schneider, *On square roots of M-matrices*, Linear Algebra Appl., 42 (1982), pp. 119–132.
2. O. Axelsson, *Iterative Solution Methods*, Cambridge University Press, Cambridge, UK, 1994.
3. S. Balay, W. Gropp, L. McInnes and B. Smith, *The Portable, Extensible Toolkit for Scientific Computing, version 2.0.13*, <http://www.mcs.anl.gov/petsc>, code and documentation, 1996.
4. M. Benzi, A. Frommer, R. Nabben and D. B. Szyld, *Algebraic theory of multiplicative Schwarz methods*, Numer. Math., 89 (2001), pp. 605–639.
5. A. Berman and R. J. Plemmons, *Nonnegative Matrices in the Mathematical Sciences*, Academic Press, New York, NY, 1979.
6. E. Bohl and I. Marek, *A model of amplification*, J. Comput. Appl. Math., 63 (1995), pp. 27–47.
7. R. Bru, F. Pedroche, and D. B. Szyld, *Additive Schwarz iterations for Markov chains*, SIAM J. Matrix Anal. Appl., 27 (2005), pp. 445–458.
8. X.-C. Cai and M. Sarkis, *A restricted additive Schwarz preconditioner for general sparse linear systems*, SIAM J. Sci. Comput., 21 (1999), pp. 792–797.
9. T. F. Chan and T. P. Mathew, *Domain Decomposition Methods*, Acta Numerica, pp. 61–143, 1994.
10. A. Frommer and D. B. Szyld, *An algebraic convergence theory for restricted additive Schwarz methods using weighted max norms*, SIAM J. Numer. Anal., 39 (2000), pp. 463–479.
11. A. Frommer and D. B. Szyld, *Weighted max norms, splittings, and overlapping additive Schwarz iterations*, Numer. Math., 83 (1999), pp. 259–278.
12. G. Karypis and V. Kumar, *MeTiS: A software package for partitioning unstructured graphs, partitioning meshes, and computing fill-reducing orderings of sparse matrices version 4.0*, University of Minnesota, Department of Computer Science and Army HPC Research Center, Minneapolis, MN, September 1998.
13. M. Luby, *A simple parallel algorithm for the maximal independent set problem*, Proc. 17th Symposium on the Theory of Computing, Association for Computing Machinery (1985), pp. 1–10.
14. I. Marek and D. B. Szyld, *Algebraic Schwarz methods for the numerical solution of Markov chains*, Linear Algebra Appl., 386 (2004), pp. 67–81.
15. J. A. Meijerink and H. A. van der Vorst, *An iterative solution method for linear systems of which the coefficient matrix is a symmetric M-matrix*, Math. Comput. 31 (1977), 148–162.
16. A. Quarteroni and A. Valli, *Domain Decomposition Methods for Partial Differential Equations*, Oxford Science Publications, Clarendon Press, Oxford, 1999.
17. Y. Saad, *Iterative Methods for Sparse Linear Systems*, Society for Industrial and Applied Mathematics, Philadelphia, PA, 2003.
18. Y. Saad and M. Schultz, *GMRES: A generalized minimum residual algorithm for solving nonsymmetric linear systems*, SIAM J. Sci. Stat. Comput., 7 (1986), pp. 856–869.
19. B. F. Smith, P. E. Bjrstad, and W. D. Gropp, *Domain Decomposition, Parallel Multilevel Methods for Elliptic Partial Differential Equations*, Cambridge University Press, Cambridge - New York - Melbourne, 1996.
20. W. J. Stewart, *Introduction to the Numerical Solution of Markov Chains*, Princeton University Press, Princeton, NJ, 1994.
21. W. J. Stewart, *MARCA Models: A collection of Markov chain models*. Available at http://www4.ncsu.edu/~billy/MARCA_Models/MARCA_Models.html.
22. A. Toselli and O. Widlund, *Domain Decomposition Methods – Algorithms and Theory*, Springer Series in Computational Mathematics, Springer-Verlag, Berlin/Heidelberg, 2005.
23. R. S. Varga, *Factorization and normalized iterative methods*, in *Boundary Problems in Differential Equations*, R.E. Langer, ed., University of Wisconsin Press, Madison, 1960.
24. R. S. Varga, *Matrix Iterative Analysis*, Prentice-Hall, Englewood Cliffs, N. J., 1962.