

**PARALLEL MONOTONE DOMAIN DECOMPOSITION
ALGORITHMS FOR NONLINEAR SINGULARLY PERTURBED
REACTION-DIFFUSION PROBLEMS OF PARABOLIC TYPE**

MATTHEW HARDY AND IGOR BOGLAEV

Mathematical Sciences Institute, Australian National University
Canberra, Australia

Institute of Fundamental Sciences, Massey University
Palmerston North, New Zealand

ABSTRACT. Recently, a monotone iterative domain decomposition algorithm has been proposed for nonlinear singularly perturbed reaction-diffusion problems of parabolic type. This paper describes a parallel implementation of the algorithm on a distributed memory cluster. Interprocess communication is effected by means of the MPI message passing library. For various domain decompositions, we give the execution time and parallel speedup on up to 16 processors. The parallel scale-up of the algorithm improves as the number of mesh points is increased.

Keywords: parabolic reaction-diffusion problem; boundary layers; implicit scheme; monotone iterative method; domain decomposition; parallel computing

1. INTRODUCTION

We are interested in the nonlinear reaction-diffusion problem of parabolic type

$$(1) \quad \begin{aligned} -\mu^2(u_{xx} + u_{yy}) + u_t &= -f(x, y, t, u), \\ (x, y, t) \in v = \omega \times (0, t_f], \quad \omega &= \{0 < x < 1, 0 < y < 1\}, \\ c^* \geq f_u \geq 0, \quad (x, y, t, u) \in \bar{v} \times (-\infty, \infty), \quad (f_u \equiv \partial f / \partial u), \end{aligned}$$

where μ is a small positive parameter and c^* is constant. The initial-boundary conditions are defined by

$$u(x, y, 0) = u^0(x, y), \quad (x, y) \in \bar{\omega}, \quad u = g, \quad (x, y, t) \in \partial\omega \times (0, t_f],$$

where $\partial\omega$ is the boundary of ω . For $\mu \ll 1$, problem (1) is singularly perturbed and characterized by boundary layers of width $\mathcal{O}(\mu |\ln \mu|)$ at $\partial\omega$. Note that the assumption $f_u \geq 0$ can always be obtained by a change of variables like $v = \exp(\lambda t)u$, where λ is a constant.

On each time level, discrete approximation of (1) leads to an algebraic system of nonlinear difference equations whose solution converges with mesh refinement to

that of the continuous problem. The algebraic system is typically solved by Newton's method, or some other iterative technique. One drawback of Newton's method is its sensitivity to the initial guess. In contrast, the method of upper and lower solutions generates a monotonically convergent sequence from any one of a wide class of initial iterates. Indeed, as shown in [1], the initial iterate may be constructed using only the boundary conditions. No knowledge of the solution is necessary to implement the algorithm.

The beowulf cluster has brought high-performance computing within reach of academe and fostered renewed interest in alternating Schwarz-type iterative algorithms. In [2] the spatial domain is partitioned into nonoverlapping boxes on which the monotone iterative method is applied at each time step. At horizontal and vertical subdomain boundaries, interfacial subdomains are introduced and corresponding linear problems generate boundary Dirichlet data for the nonoverlapping subdomains. As shown theoretically and confirmed by serial computations [2], the algorithm retains global monotonicity under such decomposition. In [5], we implemented on a parallel cluster the algorithm from [2] for an elliptic (time-independent) problem. The present paper describes a parallel implementation of the algorithm from [2] for the parabolic reaction-diffusion problem (1).

In Section 2 we define the piecewise uniform mesh on which the central difference implicit scheme converges μ -uniformly to the solution of (1). In Section 3 we describe the domain decomposition from [2] and the associated monotone iterative algorithm. We also state the main results from the convergence analysis. Section 4 discusses our parallel implementation of the algorithm. In Section 5 we give the results of numerical experiments for a model problem. For various domain decompositions we give the convergence iteration counts and execution times on up to 16 processors. Since domain decomposition enhances the algorithm's serial execution, we define parallel speedup in terms of the optimal decomposition for a given number of processors. The parallel scale-up of the algorithm improves as the mesh size is increased.

2. DISCRETE APPROXIMATION

On \bar{v} introduce the mesh $\bar{\omega}^h \times \bar{\omega}^\tau$, $\bar{\omega}^h = \bar{\omega}^{hx} \times \bar{\omega}^{hy}$ defined as

$$\begin{aligned}\bar{\omega}^{hx} &= \{x_i, 0 \leq i \leq N_x; x_0 = 0, x_{N_x} = 1; h_{xi} = x_{i+1} - x_i\}, \\ \bar{\omega}^{hy} &= \{y_j, 0 \leq j \leq N_y; y_0 = 0, y_{N_y} = 1; h_{yj} = y_{j+1} - y_j\}, \\ \bar{\omega}^\tau &= \{t_k = k\tau, 0 \leq k \leq N_\tau, N_\tau\tau = t_f\}.\end{aligned}$$

Define $\mathcal{D}_x^2 U$ and $\mathcal{D}_y^2 U$, the central difference approximations to the second x - and y -derivatives, respectively,

$$\mathcal{D}_x^2 U_{ij}^k = (h_{xi})^{-1} [(U_{i+1,j}^k - U_{ij}^k) (h_{xi})^{-1} - (U_{ij}^k - U_{i-1,j}^k) (h_{x,i-1})^{-1}],$$

$$\begin{aligned} \mathcal{D}_y^2 U_{ij}^k &= (\bar{h}_{yj})^{-1} [(U_{i,j+1}^k - U_{ij}^k) (h_{yj})^{-1} - (U_{ij}^k - U_{i,j-1}^k) (h_{y,j-1})^{-1}], \\ \bar{h}_{xi} &= 2^{-1} (h_{x,i-1} + h_{xi}), \quad \bar{h}_{yj} = 2^{-1} (h_{y,j-1} + h_{yj}), \end{aligned}$$

where $U_{ij}^k \equiv U(x_i, y_j, t_k)$.

We seek a mesh function U on $\bar{\omega}^h \times \bar{\omega}^\tau$ satisfying the implicit difference scheme

$$(2) \quad \begin{aligned} \mathcal{L}U(P, t) &= \frac{1}{\tau} U(P, t - \tau) - f(P, t, U), \quad (P, t) \in \omega^h \times \omega^\tau, \\ U(P, 0) &= u^0(P), \quad P \in \bar{\omega}^h, \quad U(P, t) = g(P, t), \quad (P, t) \in \partial\omega^h \times \omega^\tau, \end{aligned}$$

where $\mathcal{L}U = [-\mu^2 (\mathcal{D}_x^2 + \mathcal{D}_y^2) + \tau^{-1}] U$.

The meshes $\bar{\omega}^{hx}$ and $\bar{\omega}^{hy}$ are defined in the manner of [6] and are referred to as Shishkin meshes. For an arbitrary positive constant m , boundary layer thicknesses σ_x and σ_y are chosen as

$$(3) \quad \sigma_x = \min \{0.25, m\mu \ln N_x\}, \quad \sigma_y = \min \{0.25, m\mu \ln N_y\},$$

and mesh spacings $h_{x\mu}$, h_x , $h_{y\mu}$ and h_y are defined by

$$(4) \quad h_{x\mu} = \frac{4\sigma_x}{N_x}, \quad h_x = \frac{2(1 - 2\sigma_x)}{N_x}, \quad h_{y\mu} = \frac{4\sigma_y}{N_y}, \quad h_y = \frac{2(1 - 2\sigma_y)}{N_y}.$$

The mesh $\bar{\omega}^{hx}$ is constructed thus: in each of the subintervals $[0, \sigma_x]$ and $[1 - \sigma_x, 1]$ the fine mesh spacing is $h_{x\mu}$ while in the interval $[\sigma_x, 1 - \sigma_x]$ the coarse mesh spacing is h_x . The mesh $\bar{\omega}^{hy}$ is defined similarly. The difference scheme (2) on the piecewise uniform mesh (3), (4) converges μ -uniformly to the solution of the continuous problem (1);

$$\begin{aligned} \max_{t \in \bar{\omega}^\tau} \|U(t) - u(t)\|_{\bar{\omega}^h} &\leq C (d(N^{-1}) + \tau), \quad N = \min \{N_x, N_y\}, \\ \|U(t) - u(t)\|_{\bar{\omega}^h} &\equiv \max_{P \in \bar{\omega}^h} |U(P, t) - u(P, t)|, \end{aligned}$$

where d is near order 2 in N^{-1} , and constant C is independent of μ , N and τ (see [3] for details).

3. BOX-DOMAIN DECOMPOSITION ALGORITHM

For solving the nonlinear difference scheme (2), we now describe the box-domain decomposition algorithm from [2]. Consider a decomposition of ω into $M \times L$ nonoverlapping main subdomains $\omega_{m,l}$, $m = 1, \dots, M$, $l = 1, \dots, L$:

$$\omega_{m,l} = (x_{m-1}, x_m) \times (y_{l-1}, y_l), \quad x_0 = 0, x_M = 1, \quad y_0 = 0, y_L = 1.$$

Then introduce vertical interfacial subdomains θ_m , $m = 1, \dots, M - 1$:

$$\theta_m = (x_m^b, x_m^e) \times \omega^y, \quad x_m^b < x_m < x_m^e, \quad \theta_{m-1} \cap \theta_m = \emptyset,$$

and horizontal interfacial subdomains ϑ_l , $l = 1, \dots, L - 1$:

$$\vartheta_l = \omega^x \times (y_l^b, y_l^e), \quad y_l^b < y_l < y_l^e, \quad \vartheta_{l-1} \cap \vartheta_l = \emptyset.$$

With the global mesh $\bar{\omega}^h = \bar{\omega}^{hx} \times \bar{\omega}^{hy}$, we also require that

$$\{x_m^{b,e}, x_m\} \subset \omega^{hx}, \quad m = 1, \dots, M-1, \quad \{y_l^{b,e}, y_l\} \subset \omega^{hy}, \quad l = 1, \dots, L-1.$$

On each time level $t \in \omega^\tau$, the algorithm from [2] generates n_* iterates $V^{(n)}(P, t)$, $(P, t) \in \bar{\omega}^h \times \omega^\tau$ as follows.

Step 0. On the whole mesh $\bar{\omega}^h$ choose an initial mesh function $V^{(0)}(P, t)$ satisfying the boundary conditions, $V^{(0)}(P, t) = g(P, t)$, $P \in \partial\omega^h$.

For $n = 1$ to n_* do Steps 1 to 4:

Step 1. For each main subdomain $\bar{\omega}_{m,l}^h$, solve the linear difference problem

$$\begin{aligned} (\mathcal{L} + c^*)Z_{m,l}^{(n)}(P, t) &= -G^{(n-1)}(P, t), \quad P \in \omega_{m,l}^h = \omega_{m,l} \cap \omega^h, \\ G^{(n-1)}(P, t) &= \mathcal{L}V^{(n-1)}(P, t) + f(P, t, V^{(n-1)}) - \tau^{-1}V(P, t - \tau), \end{aligned}$$

with $Z_{m,l}^{(n)}(\partial\omega_{m,l}^h, t) = 0$.

Step 2. For each vertical interfacial subdomain $\bar{\theta}_m^h$, solve the linear difference problem

$$(\mathcal{L} + c^*)Z_m^{(n)}(P, t) = -G^{(n-1)}(P, t), \quad P \in \theta_m^h = \theta_m \cap \omega^h,$$

with $Z_m^{(n)}(\partial\theta_m^h, t)$ defined by the mesh functions computed in Step 1.

Step 3. For each horizontal interfacial subdomain $\bar{\vartheta}_l^h$, solve the linear difference problem

$$(\mathcal{L} + c^*)\tilde{Z}_l^{(n)}(P, t) = -G^{(n-1)}(P, t), \quad P \in \vartheta_l^h = \vartheta_l \cap \omega^h,$$

with $\tilde{Z}_l^{(n)}(\partial\vartheta_l^h, t)$ defined by the mesh functions computed in Step 2 where possible, or the mesh functions from Step 1 otherwise.

Step 4. Piece together the mesh functions from Steps 1-3 to form the new global iterate $V^{(n)}$:

$$V^{(n)}(P, t) = \begin{cases} V^{(n-1)}(P, t) + \tilde{Z}_l^{(n)}(P, t), & P \in \bar{\vartheta}_l^h, \\ V^{(n-1)}(P, t) + Z_m^{(n)}(P, t), & P \in \bar{\theta}_m^h \setminus \left\{ \bigcup_{l=1}^{L-1} \bar{\vartheta}_l^h \right\}, \\ V^{(n-1)}(P, t) + Z_{m,l}^{(n)}(P, t), & P \in \bar{\omega}_{m,l}^h \setminus \left\{ \bigcup_{m=1}^{M-1} \bar{\theta}_m^h \bigcup_{l=1}^{L-1} \bar{\vartheta}_l^h \right\}. \end{cases}$$

Step 5. Set up $V(P, t) = V^{(n_*)}(P, t)$, $P \in \bar{\omega}^h$.

Convergence of the algorithm. We say that on a time level $t \in \omega^\tau$, $\bar{V}(P, t)$ is an upper solution with respect to a given function $V(P, t - \tau)$ if it satisfies

$$\begin{aligned} \mathcal{L}\bar{V}(P, t) + f(P, t, \bar{V}) - \tau^{-1}V(P, t - \tau) &\geq 0, \quad P \in \omega^h, \\ \bar{V}(P, t) &= g(P, t), \quad P \in \partial\omega^h. \end{aligned}$$

Similarly, $\underline{V}(P, t)$ is called a lower solution with respect to $V(P, t - \tau)$ if it satisfies the reversed inequality and the boundary condition.

On each time level $t \in \omega^\tau$, we have the following convergence property of the algorithm.

Theorem 1. Let $V(P, t - \tau)$ be given and $\overline{V}^{(0)}(P, t)$ and $\underline{V}^{(0)}(P, t)$ be upper and lower solutions with respect to $V(P, t - \tau)$. From $\overline{V}^{(0)}$ and $\underline{V}^{(0)}$, the algorithm respectively generates sequences $\{\overline{V}^{(n)}\}$ and $\{\underline{V}^{(n)}\}$ which converge monotonically from above and below to the unique solution $V^*(P, t)$ of the problem

$$\mathcal{L}V(P, t) + f(P, t, V) - \tau^{-1}V(P, t - \tau) = 0, \quad P \in \omega^h,$$

$$V(P, t) = g(P, t), \quad P \in \partial\omega^h.$$

That is, at each mesh point $P \in \omega^h$, we have for $n = 0, 1, 2, \dots$

$$\underline{V}^{(n)}(P, t) \leq \underline{V}^{(n+1)}(P, t) \leq V^*(P, t) \leq \overline{V}^{(n+1)}(P, t) \leq \overline{V}^{(n)}(P, t).$$

The proof of this result may be found in [2].

Remark 1. Consider the following approach for constructing initial upper and lower solutions $\overline{V}^{(0)}(P, t)$ and $\underline{V}^{(0)}(P, t)$. Let a mesh function $R(P, r)$ be defined on $\overline{\omega}^h$ and satisfy the boundary condition $R(P, t) = g(P, t)$ on $\partial\omega^h$. Introduce the difference problems

$$\mathcal{L}Z_\nu^{(0)} = \nu |\mathcal{L}R(P, t) + f(P, t, R) - \tau^{-1}V(P, t - \tau)|, \quad P \in \omega^h,$$

$$Z_\nu^{(0)}(P, t) = 0, \quad P \in \partial\omega^h, \quad \nu = 1, -1.$$

Then the functions $\overline{V}^{(0)}(P, t) = R(P, t) + Z_1^{(0)}(P, t)$ and $\underline{V}^{(0)}(P, t) = R(P, t) + Z_{-1}^{(0)}(P, t)$ are upper and lower solutions, respectively. The proof of this result can be found in [2]. We mention that much discussions on this subject can be found in [7].

A method for constructing the initial iterate $V^{(0)}$ from an arbitrary mesh function is also given therein.

We now consider the convergence rate of the algorithm. Let h_{xm}^{b+} and h_{xm}^{b-} be the respective mesh step sizes to the right and left of x_m^b . Define h_{xm}^{e+} and h_{xm}^{e-} similarly with respect to x_m^e . Let h_{yl}^{b+} and h_{yl}^{b-} be the respective mesh step sizes above and below y_l^b . Define h_{yl}^{e+} and h_{yl}^{e-} similarly with respect to y_l^e . With averages

$$\bar{h}_{xm}^{b,e} = \frac{1}{2}(h_{xm}^{b-,e-} + h_{xm}^{b+,e+}), \quad \bar{h}_{yl}^{b,e} = \frac{1}{2}(h_{yl}^{b-,e-} + h_{yl}^{b+,e+}),$$

define parameters r^I and r^{II} , respectively associated with the x - and y -decompositions,

$$r^I = \max_{1 \leq m < M} \left\{ \frac{\mu^2}{(c^* + \tau^{-1}) \bar{h}_{xm}^b h_{xm}^{b+}}; \frac{\mu^2}{(c^* + \tau^{-1}) \bar{h}_{xm}^e h_{xm}^{e-}} \right\},$$

$$r^{II} = \max_{1 \leq l < L} \left\{ \frac{\mu^2}{(c^* + \tau^{-1}) \bar{h}_{yl}^b h_{yl}^{b+}}; \frac{\mu^2}{(c^* + \tau^{-1}) \bar{h}_{yl}^e h_{yl}^{e-}} \right\}.$$

Theorem 2. Let $U(P, t)$ be the exact solution of the nonlinear scheme (2), and at each $t \in \omega^\tau$ let $V(P, t) = V^{(n_*)}(P, t)$, $P \in \bar{\omega}^h$, where $V^{(0)}(P, t)$ is an upper or lower solution with respect to $V(P, t - \tau)$. Then we can find constant C independent of τ such that

$$(5) \quad \max_{t \in \omega^\tau} \|V(t) - U(t)\|_{\bar{\omega}^h} \leq C (c^* + \tau^{-1}) \tilde{r}^{n_*}, \quad \tilde{r} = r + r^I + r^{II},$$

here $r = c^*/(c^* + \tau^{-1})$ is the convergence rate of the undecomposed monotone iterative algorithm ($M = L = 1$).

The proof of this result is given in [2].

Remark 2. The implicit two-level difference scheme (2) is of first order with respect to τ . Since $\tilde{r} = \mathcal{O}(\tau)$, one may choose $n_* = 2$ to keep the global error of the box-domain decomposition algorithm consistent with the global error of the difference scheme (2).

Remark 3. Consider an $M \times L$ decomposition in which the mesh points are distributed equally among main subdomains. This is called a *balanced* domain decomposition. If a balanced domain decomposition has $M \geq 4$ then θ_1^h and θ_{M-1}^h overlap the boundary layers and r^I is maximal. Similarly, if a balanced decomposition has $L \geq 4$ then r^{II} is maximal. By contrast, if the interfacial subdomains are located wholly outside the boundary layers, ensuring minimality of r^I and r^{II} , then the decomposition is said to be *unbalanced*.

4. PARALLEL IMPLEMENTATION

We have implemented the box-domain decomposition algorithm on *Helix*; the distributed memory Linux cluster at Massey University, New Zealand. This comprises 65 nodes, each equipped with two Athlon MP-2100 processors, 1 GB of RAM and 2 gigabit network interface cards. The nodes are arranged on 7 24-port switches such that any two nodes are connected via at most 3 switches. Inter-processor communication is via the MPI library specification.

Because the mesh is only piecewise uniform, the linear system arising from the difference problem on a given subdomain may be nonsymmetric. Therefore, we solve all linear systems with the restarted GMRES(m) algorithm from [8]. Convergence is accelerated by the point Jacobi preconditioner. As suggested by the convergence estimate (5) and confirmed by the serial computations of balanced and unbalanced domain decompositions in [2], the latter require fewer iterations for convergence. However, in order to balance the computational load of the main subdomains, one would need to divide the larger boundary layer subdomains across several processors. Thus, the algorithmic advantage would be at least partially offset by the need for inter-processor communication during the solution of the main subdomain linear problems.

Therefore, our implementation is at the first level. That is, each subdomain is wholly assigned to a processor and we do not parallelize GMRES. We henceforth consider only balanced domain decompositions.

Suppose we have an $M \times L$ decomposition and P processors. Assuming that $M \geq 4$ and $L \geq 4$, there are $ML/16$ main subdomains in each of the four corners $[0, \sigma_x] \times [0, \sigma_y]$, $[1 - \sigma_x, 1] \times [0, \sigma_y]$, $[0, \sigma_x] \times [1 - \sigma_y, 1]$ and $[1 - \sigma_x, 1] \times [1 - \sigma_y, 1]$. Thus, there are $ML/4$ main subdomains in which the respective x - and y -mesh spacings are $h_{x\mu}$ and $h_{y\mu}$. Let us denote this class by F-F (fine-fine). Next, in each of the regions $[\sigma_x, 1 - \sigma_x] \times [0, \sigma_y]$, $[\sigma_x, 1 - \sigma_x] \times [1 - \sigma_y, 1]$, $[0, \sigma_x] \times [\sigma_y, 1 - \sigma_y]$ and $[1 - \sigma_x, 1] \times [\sigma_y, 1 - \sigma_y]$, there are $ML/8$ main subdomains in which the respective x - and y - mesh spacings are either $h_{x\mu}$ and h_y or h_x and $h_{y\mu}$. In the numerical experiments of the next section we take $N_x=N_y$ and hence these $ML/2$ subdomains are equivalent. We label this class F-C (fine-coarse). Finally, the region $[\sigma_x, 1 - \sigma_x] \times [\sigma_y, 1 - \sigma_y]$ is covered by $ML/4$ main subdomains in which the respective x - and y -mesh spacings are h_x and h_y . We denote this class C-C (coarse-coarse).

The linear difference problem on $\omega_{m,l}^h$ may be written in the form

$$[-\mu^2(\mathcal{D}_x^2 + \mathcal{D}_y^2) + \tau^{-1} + c^*] Z_{m,l}^{(n)} = -G^{(n-1)}(P, t), P \in \omega_{m,l}^h.$$

If the mesh $\omega_{m,l}^h$ is uniform in each direction, then the rows of the coefficient matrix $-\mu^2(\mathcal{D}_x^2 + \mathcal{D}_y^2)$ have entries which sum to zero (excepting those rows which correspond to mesh points adjacent to $\partial\omega_{m,l}^h$). Strict diagonal dominance of the coefficient matrix in the linear difference problem is assured by the positive parameters τ^{-1} and c^* and the relative contribution of these parameters increases with the subdomain mesh spacing. Therefore, the GMRES workload is least for the C-C subdomains and greatest for the F-F subdomains.

In order to balance the load of Step 1 of the algorithm, we divide each class equally among the processors. If P divides $ML/4$, each processor is assigned $ML/(4P)$ main subdomains of class F-F, $ML/(2P)$ main subdomains of class F-C and $ML/(4P)$ main subdomains of class C-C. If P does not divide $ML/4$ then load balancing requires second level parallelization and we do not implement the $M \times L$ decomposition on P processors. Similar considerations apply to decompositions in which M or L is equal to one, except that there are at most two classes of main subdomain. The vertical interfacial subdomains are shared as equally as possible among the P processors $\{0, 1, \dots, P - 1\}$: $\bar{\theta}_m^h \mapsto \text{mod}(m - 1, P)$. The assignment of horizontal interfacial subdomains is similar: $\bar{\vartheta}_l^h \mapsto \text{mod}(l - 1, P)$.

Consider the distribution of an 8×4 decomposition on four processors, shown schematically in Figure 1. During Step 1 of the algorithm, each processor solves over its assigned main subdomains and this workload is balanced. During Step 2, each processor solves over its assigned vertical interfacial subdomains. In contrast to the

other processors which each have two, Processor 3 has only one vertical interfacial subdomain and so is idle for approximately half of Step 2. In Step 3, Processors 0,1,2 each solve over their assigned horizontal interfacial subdomain while Processor 3 remains idle. This idle time results in a loss of computational efficiency.

Another overhead of any parallel implementation is that of inter-processor communication. Before each of the algorithm’s three steps, data must be transferred between subdomains. In Figure 1 we have indicated the data which must be transferred from Processor 0 to Processor 1 before the latter can solve over its main subdomains. From the left-most vertical interfacial subdomain on Processor 0, we must transfer the three blocks of data **a**, **b** and **c** to main subdomains on Processor 1. From the horizontal interfacial subdomain on Processor 0, we must transfer the three blocks of data **d**, **e** and **f** to Processor 1. In order to minimize the inter-processor communication, the data blocks **a**–**f** are buffered and sent as one message. This saving in communication comes at the relatively small (local) cost of composing the message on Processor 0 and decomposing the received message on Processor 1.




FIGURE 1. A schematic of an 8×4 decomposition and the assignment of the subdomains to four processors $\{0, 1, 2, 3\}$. The mesh spacing classes of the main subdomains are also indicated. The six blocks of data **a**–**f** must be transferred from Processor 0 to Processor 1 before Step 1 of the algorithm.

5. NUMERICAL EXPERIMENTS

We now apply the algorithm to the reaction-diffusion problem

$$-\mu^2(u_{xx} + u_{yy}) + u_t = -\frac{u - 4}{5 - u},$$

$$(x, y, t) \in \omega \times (0, 1], \quad \omega = \{0 < x < 1, 0 < y < 1\},$$

$$u(\omega, 0) = 0, \quad u(\partial\omega, 0) = 1, \quad u(\partial\omega, t) = 1, \quad t \in (0, 1],$$

which models the biological Michaelis-Menten process without inhibition [4]. The steady state solution to the reduced problem ($\mu = 0$) is $u_r = 4$. For $\mu \ll 1$ the problem is singularly perturbed and the steady state solution increases sharply from $u = 1$ on $\partial\omega$ to $u = 4$ on the interior. The solution to the parabolic problem approaches this steady state with time.

For the test problem we solve the nonlinear difference scheme (2) by the domain decomposition algorithm. With the mesh (3), (4), we take $N_x = N_y = N$. We fix the perturbation parameter $\mu = 10^{-3}$, take $N = 256, 512$ or 1024 and $\tau = 0.1$ or 0.05 . Given up to 16 processors, we are interested in the execution time of the algorithm under various domain decompositions. We suppose that $\{M, L\} \subset \{1, 4, 8, 16, 32\}$ and that the interfacial subdomains are chosen to be either all minimal or all maximal. For the minimal interfacial subdomains the number of mesh points in the x - and y -directions is three, and for the maximal interfacial subdomains the numbers of mesh points in the x - and y -directions are $N/M - 2$ and $N/L - 2$, respectively, (it means that the maximal interfacial subdomains do not overlap each other).

The numerical solution at $t_0 = 0$ is simply given by the initial-boundary conditions $V(\omega^h, t_0) = 0, V(\partial\omega^h, t_0) = 1$. The mesh function $\underline{V}^{(0)}(P, t_1)$ defined by $\underline{V}^{(0)}(P, t_1) = V(P, t_0), P \in \bar{\omega}^h$ is clearly a lower solution with respect to $V(P, t_0)$. We initiate the algorithm with $\underline{V}^{(0)}(P, t_1)$ and thus generate a sequence of lower solutions. At each time level t_k , we define a converged solution $V(P, t_k) = \underline{V}^{(n_*)}(P, t_k)$ with $n_* = n_*(t_k)$ minimal subject to $\|\underline{V}^{(n_*)}(t_k) - \underline{V}^{(n_*-1)}(t_k)\|_{\bar{\omega}^h} < \delta$, where δ is a specified tolerance. At the next time level, t_{k+1} , we require an initial iterate that is a lower solution with respect to $V(P, t_k)$. Since the boundary condition and function $f(u) = (u - 4)/(5 - u)$ are independent of time, we may choose $\underline{V}^{(0)}(P, t_{k+1}) = V(P, t_k), P \in \bar{\omega}^h$. Indeed, since $\underline{V}^{(0)}(P, t_1) = V(P, t_0) \leq \underline{V}(P, t_1)$, on time level t_2 with $\underline{V}^{(0)}(P, t_2) = \underline{V}(P, t_1)$, we have

$$\begin{aligned} \mathcal{L}\underline{V}^{(0)}(P, t_2) + f(P, \underline{V}^{(0)}(P, t_2)) - \tau^{-1}\underline{V}(P, t_1) &\leq \\ \mathcal{L}\underline{V}(P, t_1) + f(P, \underline{V}(P, t_1)) - \tau^{-1}V(P, t_0) &\leq 0, \end{aligned}$$

where $P \in \bar{\omega}^h$ and $V^{(0)}(P, t_2) = g(P), P \in \partial\omega^h$. Thus, $\underline{V}^{(0)}(P, t_2) = \underline{V}(P, t_1)$ is a lower solution, and by induction on k , we prove the required result. Now, from Theorem 1, it follows by induction on k that the mesh function $\bar{V}(P, t_{k+1})$ defined by $\bar{V}(\omega^h, t_{k+1}) = 4, \bar{V}(\partial\omega^h, t_{k+1}) = 1$ is an upper solution with respect to $V(P, t_k)$ and thus our computed mesh functions satisfy

$$0 \leq \underline{V}^{(n)}(P, t_k) \leq 4, \quad P \in \bar{\omega}^h, \quad 0 \leq n \leq n_*, \quad 0 \leq k \leq N_\tau,$$

and hence we may suppose that $f_u = 1/(5 - u)^2$ is bounded below and above by $c_* = 1/25$ and $c^* = 1$, respectively. We take as our convergence tolerance $\delta = 10^{-5}$.

We list in Table 1 average convergence iteration counts for the $M \times L$ domain decomposition algorithm. Each major cell of the table corresponds to a choice of N and τ and within each major cell we give the results corresponding to twenty-five $M \times L$ decompositions with minimal interfacial subdomains. We mention that domain decomposition does not increase the iteration count if the interfacial subdomains are chosen maximally. Because the final simulation time is $t = 1$, the average is over 10 time steps for $\tau = 0.1$ and 20 time steps for $\tau = 0.05$. For a given decomposition, the average iteration count is smaller for $\tau = 0.05$ than for $\tau = 0.1$. This simply reflects the smaller change in solution between time steps.

	τ	0.10					0.05				
N	$L \setminus M$	1	4	8	16	32	1	4	8	16	32
256	1	5.00	6.00	7.60	7.60	7.70	4.00	5.00	6.00	6.00	6.00
	4	6.00	6.00	7.60	7.60	8.00	5.00	5.00	6.00	6.00	6.00
	8	7.60	7.60	9.00	9.00	9.00	6.00	6.00	6.00	6.00	6.00
	16	7.60	7.60	9.00	9.00	9.00	6.00	6.00	6.00	6.00	6.00
	32	7.70	7.70	9.00	9.00	9.00	6.00	6.00	6.00	6.00	6.00
512	1	5.00	7.00	11.00	11.00	11.00	4.05	5.65	8.00	8.00	8.00
	4	7.00	7.00	11.00	11.00	11.00	5.65	6.00	8.00	8.00	8.00
	8	11.00	11.00	13.00	13.00	13.00	8.00	8.00	9.00	9.00	9.00
	16	11.00	11.00	13.00	13.00	13.00	8.00	8.00	9.00	9.00	9.00
	32	11.00	11.00	13.00	13.00	13.00	8.00	8.00	9.00	9.00	9.00
1024	1	5.00	10.00	18.10	18.10	18.10	4.05	7.00	12.20	12.20	12.20
	4	10.00	10.00	18.10	18.10	18.10	7.00	7.00	12.20	12.20	12.20
	8	18.10	18.10	23.00	23.00	23.00	12.20	12.20	15.00	15.00	15.00
	16	18.10	18.10	23.00	23.00	23.00	12.20	12.20	15.00	15.00	15.00
	32	18.10	18.10	23.00	23.00	23.00	12.20	12.20	15.00	15.00	15.00

TABLE 1. The average convergence iteration count for various τ , N and $M \times L$ decompositions with minimal interfacial subdomains.

In Tables 2 and 3 we give the execution times of the algorithm for $\tau = 0.1$ and $\tau = 0.05$, respectively. Where there is some choice for the interfacial subdomains, results corresponding to minimal and maximal interfacial subdomains are given above and below the line, respectively.

Consider first the serial implementation corresponding to $P = 1$. For $N = 256$ and $\tau = 0.1$, the execution time of the undecomposed algorithm ($M = 1$, $L = 1$) is 19.59 seconds. If one chooses the interfacial subdomains minimally, then domain decomposition reduces the serial execution time. The execution time is minimized to 4.77 seconds by the 4×32 decomposition with minimal interfacial subdomains. For

each τ , it is evident from Tables 2 and 3 that the algorithm's serial execution is fastest on the 32×32 decomposition with minimal interfacial subdomains. Clearly, we have a notion of *serial speedup*, which derives purely from the domain decomposition. For $P \geq 2$ our parallel implementation exploits the natural parallelism of the algorithm at Steps 1-3. Consider now the problem with $N = 512$ and $\tau = 0.1$. For three $M \times L$ decompositions and each value of P , we give in Table 4 the execution time, number of inter-process subdomain data transfers per global iteration (Steps 1 through 4 of the algorithm) and the total number of MPI messages per global iteration by which this data is transmitted. The serial code ($P = 1$) executes fastest on the 32×32 decomposition. For $P \leq 4$, the number of MPI communications is the same for all three listed decompositions. Therefore, we see that the 32×32 decomposition allows for the fastest execution on $P \leq 4$ processors. For $P \geq 8$ we find that the number of MPI communications necessitated by the 32×8 or 8×32 decomposition is significantly less than that by the 32×32 decomposition. As a result, we observe that the 32×32 decomposition does not minimize the execution time when $P \geq 8$. On the other hand, if N increases to 1024, the linear algebra arising from the 32×32 decomposition becomes increasingly significant and the communication overhead is increasingly justified. Thus, we see from Tables 2 and 3 that the 32×32 , minimal interfacial subdomain decomposition minimizes the execution time for any number of processors when $N = 1024$.

P	32×32			32×8			8×32		
1	54.16	0	0	62.16	0	0	59.88	0	0
2	28.08	4928	6	31.41	1160	6	31.08	1160	6
4	15.58	6400	36	16.44	1740	36	17.26	1620	36
8	9.29	7136	168	8.93	1968	161	9.36	1786	161
16	6.88	9008	720	7.01	2212	489	6.26	2118	502

TABLE 4. For three $M \times L$ decompositions and five values of P , a triple giving the execution time, number of inter-process subdomain transfers per global iteration and number of MPI messages per global iteration. The problem has $N = 512$ and $\tau = 0.1$.

For each combination of τ , N , and P , we define the parallel speedup of our implementation with reference to the optimal domain decomposition. Let $T(\tau, N, P)$ be the minimal execution time over all domain decompositions on P processors. Thus, $T(\tau, N, P)$ is the smallest number in the major cell of Tables 2 and 3 corresponding to τ , N and P . We define the parallel speedup on P processors as the ratio

$$S(\tau, N, P) = \frac{T(\tau, N, 1)}{T(\tau, N, P)}.$$

In Figure 2 we show the parallel speedup as a function of P for $\tau = 0.1$ and the three values of N . The corresponding graph for $\tau = 0.05$ is shown in Figure 3. It is evident from these graphs that the parallel scale-up of the algorithm improves as the number of mesh points is increased.




FIGURE 2. The parallel speedup $S(\tau = 0.1, N, P)$ as a function of number of processors P for three values of N and $\tau = 0.1$.




FIGURE 3. The parallel speedup $S(\tau = 0.05, N, P)$ as a function of number of processors P for three values of N and $\tau = 0.05$.

6. CONCLUSION

We have described the domain decomposition algorithm from [2] for singularly perturbed reaction-diffusion problems of parabolic type. A parallel implementation of the algorithm on a distributed memory cluster has been discussed.

From numerical experiments with a model test problem, we have found that if one chooses the interfacial subdomains maximally then the convergence iteration count remains at the undecomposed value. However, our measurements of wall time have shown that the algorithm executes fastest when the interfacial subdomains are chosen minimally. We have observed that domain decomposition can reduce the serial execution time of the algorithm. We have thus measured parallel speedup on P processors with reference to the best decomposition for P processors. The parallel scale-up of the algorithm improves as the number of mesh points is increased. With $N = 1024$ and $\tau = 0.1$ or 0.05 , the parallel speedup on 16 processors is 13.6, corresponding to a parallel efficiency of 85.0

If one choose the interfacial subdomains minimally, then the domain decomposition algorithm reduces the serial execution time of the undecomposed algorithm. The parallel domain decomposition algorithm executes fastest with the minimal interfacial subdomains. When the number of mesh points increases, then the 32×32 , minimal interfacial subdomain decomposition minimizes the parallel execution time for any number of processors.

REFERENCES

- [1] I. Boglaev, Monotone iterative algorithms for a nonlinear singularly perturbed parabolic problem, *J. Comput. Appl. Math.* 172, 313–335 (2004).
- [2] I. Boglaev, M.P. Hardy, Monotone box-domain decomposition algorithms for nonlinear singularly perturbed reaction-diffusion problems, *Adv. Difference Equ.* Vol. 2006, Article ID 70325, 38 pages (2006)// DOI:10.1155/ADE/2006/70325.
- [3] I. Boglaev, M.P. Hardy, Uniform convergence of a weighted average scheme for a nonlinear reaction-diffusion problem, *J. Comput. Appl. Math.* 200, 705–721 (2007).
- [4] E. Bohl, *Finite Modelle gewöhnlicher Randwertaufgaben*, Teubner, Stuttgart (1981).
- [5] M.P. Hardy, I. Boglaev, Parallel implementation of a monotone domain decomposition algorithm for nonlinear reaction-diffusion problems, *ANZIAM J.* 46, C290–C303 (2005).
- [6] J.J.H. Miller, E. O’Riordan, G.I. Shishkin, *Fitted Numerical Methods for Singular Perturbation Problems*, World Scientific, Singapore (1996).
- [7] C.V. Pao, Finite difference reaction diffusion equations with nonlinear boundary conditions, *Numer. Meth. Part. Diff. Eqs.* 11, 355–374 (1995).
- [8] Y. Saad, M.H. Schultz, A generalized minimal residual method for solving nonsymmetric linear systems, *SIAM J. Sci. Stat. Comput.* 7, 856–869 (1986).