

ASYNCHRONOUS DOMAIN DECOMPOSITION METHODS FOR NONLINEAR PDES*

FAYÇAL CHAOUQUI[†], EDMOND CHOW[‡], AND DANIEL B. SZYLD[†]

Abstract. One- and two-level parallel asynchronous methods for the numerical solution of nonlinear systems of equations, especially those arising from (nonlinear) partial differential equations, are studied. The proposed methods are based on domain decomposition techniques. Local convergence theorems are presented in several cases, with appropriate hypotheses. Computational results on a shared memory multiprocessor machine for various problems exhibiting nonlinearities are reported, illustrating the potential of these asynchronous methods, especially for heterogeneous clusters.

Key words. asynchronous iterations, nonlinear problems, domain decomposition, partial differential equations, two-level methods

AMS subject classifications. 65N22, 65M55, 65F10, 65F50,

1. Introduction. Asynchronous iterative methods are currently undergoing a resurgence in popularity due to the dramatic increase of parallelism in modern computers. The good performance of asynchronous methods is due to not needing to synchronize the computational tasks, hence minimizing idle time, i.e., time for which some processors are inactive but can be used. The effect of data exchange in the asynchronous method is then less pronounced; see, e.g., [1, 3, 5, 22]. In these methods, the iterations (or updates) are carried out in parallel by processors in an arbitrary order and without any synchronization. These asynchronous iterative methods are especially attractive when the network of processors is heterogeneous, or the communication costs are large, or when different processors have different loads, e.g., when each processor solves local problems with different physical properties or with different types of nonlinearities.

In this paper, we study the parallel solution of nonlinear systems of equations. These often arise from the finite element or finite difference discretization of nonlinear differential equations such as nonlinear convection-diffusion problems. We present new one- and two-level algorithms based on domain decomposition methods. The domain of the differential equation is subdivided into several possibly overlapping subdomains, and the computational tasks are assigned to different processors whereby the local components of the iterate vector can be updated without any order nor synchronization.

Our point of departure is the restricted additive Schwarz (RAS) iterative method for linear systems [11], which we review in Section 2. We then introduce a RAS method for nonlinear problems and analyze its local convergence properties. We note that RAS has been used as preconditioner for the linear systems arising in Newton’s methods for nonlinear problems; see, e.g., [9, 10, 17, 29]. In these methods, the linear systems at each step are solved using a Krylov subspace method preconditioned with RAS. Since we have in mind parallel computers with high communication costs, the use of Krylov methods may not be suitable, since each iteration usually entails orthogonalizations with the concomitant inner products, which in turn entail communication among all processors. This is the reason why we concentrate

*Received August 15, 2021. Accepted September 15, 2022. Published online on November 2, 2022. Recommended by O. Widlund. This research is supported in part by the U.S. Department of Energy under grant DE-SC0016578.

[†]Department of Mathematics, Temple University, Philadelphia, PA, USA
(fchaouqui, szyld}@temple.edu).

[‡]School of Computational Science and Engineering, Georgia Institute of Technology, Atlanta, GA, USA
(echow@cc.gatech.edu).

on RAS as a solver, not as a preconditioner. We note that earlier work has been done on using RAS as a preconditioner; cf. [12, 17]. We refer the reader also to the recent work [28] for two-level RAS preconditioners using additive/multiplicative coupling of coarse spaces, to [13] for substructured coarse correction, and to [24] for asynchronous multiplicative coarse correction.

Our contribution includes the introduction of the asynchronous version of the nonlinear RAS method in Section 4, where we also state convergence results and discuss its implementation. We further introduce a coarse grid correction in Section 5 and discuss convergence results of the two-level method. Throughout the paper, we illustrate the performance of the methods with numerical experiments.

2. Domain decomposition methods. Our goal is to study the parallel solution of discretized general nonlinear elliptic equations in \mathbb{R}^d ($d \geq 2$) of the form

$$(2.1) \quad L(u) := - \sum_{i=1}^d \frac{\partial}{\partial x_i} a_i(x, u(x), \nabla u(x)) + a_0(x, u(x), \nabla u(x)) = f(x), \quad \text{in } \Omega,$$

with homogeneous Dirichlet boundary conditions, i.e., $u = 0$ on $\partial\Omega$. Here Ω is a bounded region in \mathbb{R}^d with Lipschitz continuous boundary. Let $p = (p_0, p_1, \dots, p_d) \in \mathbb{R}^{d+1}$ be such that one can define the functions a_i as $a_i(x, u(x), \nabla u(x)) = a_i(x, p)$, for $i = 0, \dots, d$. Existence and uniqueness of a solution to problem (2.1) can be stated under certain hypotheses on the coefficients $a_i(x, p)$ as follows. Let $c, C > 0$ be such that for $i = 0, \dots, d$, $a_i(x, p)$ satisfy

- $a_i \in C^1(\Omega \times \mathbb{R}^{d+1})$,
- $\max \left\{ |a_i|, \left| \frac{\partial a_i}{\partial x_j} \right|, \left| \frac{\partial a_i}{\partial p_k} \right| \right\} \leq C$, for $i, k = 0, \dots, d, j = 1, \dots, d$, and
- $\sum_{i,j=0}^d \frac{\partial a_i(x,p)}{\partial p_j} \xi_i \xi_j \geq c \sum_{i=0}^d \xi_i^2$, for $(\xi_0, \dots, \xi_d) \in \mathbb{R}^{d+1}$.

Then, problem (2.1) is well posed in $H^1(\Omega)$, where $H^1(\Omega)$ is the space of functions $v \in L^2(\Omega)$ such that $\nabla v \in L^2(\Omega)$, equipped with the norm $\|v\|_{H^1(\Omega)}^2 := \|v\|_{L^2(\Omega)}^2 + \|\nabla v\|_{L^2(\Omega)}^2$; see, e.g., [31].

After discretization (e.g., with finite differences or finite elements), problem (2.1) yields a large nonlinear system of equations of the form

$$(2.2) \quad F(u) = B(u) - f = 0,$$

where $B, F : V \mapsto V$, $V := \mathbb{R}^n$, and n is the total number of unknowns. Here, we abuse the notation and keep the same symbol for the continuous and discrete variables. Throughout this article, we denote by $u^* \in \mathbb{R}^n$ the vector that represents the solution of (2.2).

2.1. Restricted additive Schwarz (RAS) for linear systems. We introduce in the next section the restricted additive Schwarz (RAS) method for nonlinear equations. We begin here by reviewing RAS for the case that $F(u)$ is linear. Let us thus consider

$$F(u) = Au - b,$$

where $A \in \mathbb{R}^{n \times n}$ and $b \in \mathbb{R}^n$. We assume that F has a simple zero u^* , i.e., that A is nonsingular, or in other words, A is of full rank, i.e., $\text{rank } A = n$. The RAS introduced in [11] provides a means of constructing parallel iterative solvers based on domain decomposition techniques. The general philosophy of RAS is that the local problems (i.e., in each subdomain) are solved with overlap, but the variables communicated to the other local problems are those corresponding to subdomains without overlap.

Let $S = \{1, \dots, n\}$ be the set of indices of all variables, and let the sets $S_i \subset S$, $i = 1, \dots, P$, form a non-overlapping partitioning, i.e.,

$$\bigcup_{i=1}^P S_i = S, \quad S_i \cap S_j = \emptyset, \quad i \neq j, \in \{1, \dots, P\}.$$

These correspond to the variables of the subdomains without overlap. For the overlapping subdomains, for any $\delta > 0$, we define $S_i \subset S_{i,\delta} \subset S$, $i = 1, \dots, P$, that is a covering of S . We set $V_{i,\delta} := \mathbb{R}^{n_{i,\delta}}$, with $n_{i,\delta} = |S_{i,\delta}|$ the cardinality of the set of variables corresponding to the i th subdomain with overlap δ . We define also the restriction mappings $R_{i,\delta}: V \mapsto V_{i,\delta}$, $i = 1, \dots, P$. These restriction operators satisfy the identities

$$\begin{aligned} R_{i,\delta} R_{i,\delta}^\top &= I_{V_{i,\delta}}, \\ R_{i,\delta}^\top R_{i,\delta} |_{R_{i,\delta}^\top(V_{i,\delta})} &= I_{|R_{i,\delta}^\top(V_{i,\delta})}. \end{aligned}$$

Moreover, we have

$$\sum_{i=1}^P R_{i,0}^\top R_{i,0} = I,$$

which is known as the partition of unity condition.

We are ready to define the RAS iterative method. Given an initial vector u^0 , the iteration is given by

$$(2.3) \quad u^{k+1} = u^k + \sum_{i=1}^P R_{i,0}^\top A_{i,\delta}^{-1} R_{i,\delta} (b - Au^k),$$

where $A_{i,\delta} := R_{i,\delta} A R_{i,\delta}^\top: V_{i,\delta} \mapsto V_{i,\delta}$ is the coefficient matrix for the local problem. Observe that in (2.3) the local problem is solved with the overlap, but only the variables without the overlap contribute to the next iterate. The iteration operator corresponding to (2.3) is thus given by

$$(2.4) \quad T_{A,\delta} = I - \sum_{i=1}^P R_{i,0}^\top A_{i,\delta}^{-1} R_{i,\delta} A.$$

We present a convergence and comparison result for (linear) RAS from [23]. We need some notation first. Let \geq ($>$) in \mathbb{R}^n and $\mathbb{R}^{n \times n}$ denote the natural elementwise partial ordering, i.e., for $x, y \in \mathbb{R}^n$, $x \geq y$ ($x > y$) if $x_j \geq y_j$ ($x_j > y_j$), for $j = 1, \dots, n$. Let $w \in \mathbb{R}^n$, $w > 0$, and let $\|\cdot\|_w$ denote the weighted max-norm operator corresponding to the vector norm $\|x\|_w = \max_{j=1,\dots,n} |x_j/w_j|$. A nonsingular matrix A is called an M -matrix if its off-diagonal elements are nonpositive and the inverse is nonnegative, i.e., $A^{-1} \geq 0$.

THEOREM 2.1 ([23]). *Let A be a nonsingular M -matrix. Let $w > 0$ such that $Aw > 0$. Then, if $\delta \geq \delta'$,*

$$\|T_{A,\delta}\|_w \leq \|T_{A,\delta'}\|_w < 1.$$

Theorem 2.1 shows that the (linear) RAS iterative method case converges linearly, and that the larger the overlap, the faster the asymptotic convergence. Recall that the classical additive Schwarz method may fail to converge as an iterative method. The main difference is that here

the variables of the overlap are not included in the new iterate (and thus there is no “double counting”); see [18, 23].

For completeness, we remark that when RAS is used as a preconditioner for Krylov subspace methods, it follows from (2.4) that the action of the RAS preconditioner M^{-1} on the global matrix A is given by

$$M^{-1}A = \sum_{i=1}^P R_{i,0}^\top A_{i,\delta}^{-1} R_{i,\delta} A.$$

As mentioned in the introduction, in this paper we concentrate on asynchronous parallel solvers, and therefore we do not consider Krylov subspace methods (or preconditioning), since the orthogonalization processes lead to synchronization points.

2.2. Nonlinear RAS (NLRAS). We return to the case of nonlinear F and now describe the nonlinear iterative RAS (NLRAS) algorithm. While this algorithm is new, we base our presentation on [16], where general Schwarz methods for nonlinear problems are discussed. Define the restriction operators

$$\begin{aligned} F_{i,\delta} : V &\mapsto V_{i,\delta} \\ v &\mapsto R_{i,\delta} F(v), \end{aligned}$$

for $i = 1, \dots, P$. For any given $u = u^k \in \mathbb{R}^d$, define $v_{i,\delta}$ as the solution of the local nonlinear problem

$$(2.5) \quad F_{i,\delta}(u + R_{i,\delta}^\top v_{i,\delta}) = 0, \quad \text{for } i = 1, \dots, P.$$

We can then define the NLRAS fixed point iteration as

$$(2.6) \quad u^{k+1} = u^k + \sum_{i=1}^P R_{i,0}^\top v_{i,\delta}^k.$$

Note here again the philosophy of RAS: we solve the local (nonlinear) problem (2.5) with the overlap, but only consider the non-overlapping variables as a contribution to the next iterate.

We provide now a result showing that (2.6) converges locally and has an asymptotic convergence rate similar to that of the linear case. To this end, let T define the nonlinear mapping

$$(2.7) \quad T(u) = u + \sum_{i=1}^P R_{i,0}^\top v_{i,\delta},$$

where $v_{i,\delta}$ is the solution of the local nonlinear problem (2.5), and recall that by u^* we denote the solution of (2.2). We shall prove that T is a contraction. Hence it admits a fixed point and the iterative process (2.6) converges to u^* ; cf. [14].

THEOREM 2.2. *Suppose that $F'(u^*)$ is a nonsingular M -matrix. Let $w > 0$ be such that $F'(u^*)w > 0$. Then, there exist $0 < \zeta < 1$, and \mathcal{U} a neighborhood of u^* so that*

$$\|T(u) - T(u')\|_w \leq \zeta \|u - u'\|_w \quad \text{for all } u, u' \in \mathcal{U},$$

i.e., T has a fixed point u^ and it is the solution of (2.2).*

Proof. By definition, we have

$$\begin{aligned}
 (2.8) \quad T(u) - T(u') &= \left(u + \sum_{i=1}^P R_{i,0}^\top v_{i,\delta} \right) - \left(u' + \sum_{i=1}^P R_{i,0}^\top v'_{i,\delta} \right) \\
 &= u - u' + \sum_{i=1}^P R_{i,0}^\top (v_{i,\delta} - v'_{i,\delta}).
 \end{aligned}$$

Let us define the multidimensional differential quotient of F by

$$DF(v, v') = \int_0^1 F'(v + t(v' - v)) dt, \quad v, v' \in \mathbb{R}^n.$$

Let us denote the matrix $J_{i,\delta} := DF(u + R_{i,\delta}^\top v_{i,\delta}, u' + R_{i,\delta}^\top v'_{i,\delta})$. We have that $v_{i,\delta}$ and $v'_{i,\delta}$ satisfy

$$(2.9) \quad v_{i,\delta} - v'_{i,\delta} = -(R_{i,\delta} J_{i,\delta} R_{i,\delta}^\top)^{-1} R_{i,\delta} J_{i,\delta} (u - u').$$

Indeed, using (2.5), we have that $v_{i,\delta}$ and $v'_{i,\delta}$ satisfy

$$\begin{aligned}
 0 &= F_{i,\delta}(u + R_{i,\delta}^\top v_{i,\delta}) - F_{i,\delta}(u' + R_{i,\delta}^\top v'_{i,\delta}), \\
 0 &= R_{i,\delta}(F(u + R_{i,\delta}^\top v_{i,\delta}) - F(u' + R_{i,\delta}^\top v'_{i,\delta})), \\
 0 &= R_{i,\delta} DF(u + R_{i,\delta}^\top v_{i,\delta}, u' + R_{i,\delta}^\top v'_{i,\delta})(u - u' + R_{i,\delta}^\top (v_{i,\delta} - v'_{i,\delta})), \\
 0 &= R_{i,\delta} J_{i,\delta} (u - u' + R_{i,\delta}^\top (v_{i,\delta} - v'_{i,\delta})),
 \end{aligned}$$

from which we obtain

$$R_{i,\delta} J_{i,\delta} (u - u') = -R_{i,\delta} J_{i,\delta} R_{i,\delta}^\top (v_{i,\delta} - v'_{i,\delta}).$$

Hence, substituting (2.9) in (2.8), we obtain

$$\begin{aligned}
 (2.10) \quad T(u) - T(u') &= u - u' - \sum_{i=1}^P R_{i,0}^\top (R_{i,\delta} J_{i,\delta} R_{i,\delta}^\top)^{-1} R_{i,\delta} J_{i,\delta} (u - u') \\
 &= \left(I - \sum_{i=1}^P R_{i,0}^\top (R_{i,\delta} J_{i,\delta} R_{i,\delta}^\top)^{-1} R_{i,\delta} J_{i,\delta} \right) (u - u').
 \end{aligned}$$

To complete the proof, we remark that by the definition of DF ,

$$(R_{i,\delta} DF(u, u') R_{i,\delta}^\top)^{-1} R_{i,\delta} DF(u, u') \rightarrow (R_{i,\delta} F'(u^*) R_{i,\delta}^\top)^{-1} R_{i,\delta} F'(u^*) \quad \text{as } u, u' \rightarrow u^*.$$

Moreover, formula (2.9) shows that $v_{i,\delta}, v'_{i,\delta} \rightarrow 0$, when $u, u' \rightarrow u^*$. From this it follows that

$$I - \sum_{i=1}^P R_{i,0}^\top (R_{i,\delta} J_{i,\delta} R_{i,\delta}^\top)^{-1} R_{i,\delta} J_{i,\delta} \rightarrow T_{F'(u^*),\delta} \quad \text{as } u, u' \rightarrow u^*,$$

where we used the notation in (2.4). This shows that for $\|T_{F'(u^*),\delta}\|_w < \zeta < 1$, there exists a neighborhood \mathcal{U} of u^* such that

$$\left\| I - \sum_{i=1}^P R_{i,0}^\top (R_{i,\delta} J_{i,\delta} R_{i,\delta}^\top)^{-1} R_{i,\delta} J_{i,\delta} \right\|_w \leq \zeta < 1.$$

It suffices then to take the norm $\|\cdot\|_w$ on both sides of (2.10) to obtain the stated result. \square

Theorem 2.2 shows that the iterative procedure $u^{k+1} = T(u^k)$ for $u^0 \in \mathcal{U}$ converges to $u = T(u)$ for some u . This implies that the NLRAS converges to the point $u = u^*$, the unique solution of (2.2). Moreover, the same result shows that the mapping T in (2.7) is differentiable at $u = u^*$, with $T'(u^*) = T_{F'(u^*),\delta}$. This also demonstrates that NLRAS with an overlap $\delta' \geq \delta$ has in general a faster local convergence rate since, by Theorem 2.1, $\|T_{F'(u^*),\delta'}\|_w \leq \|T_{F'(u^*),\delta}\|_w$.

We remark that the hypothesis that $F'(u^*)$ is a nonsingular M -matrix may or may not be satisfied for many practical problems. Although we can only prove convergence in this case, numerical experiments indicate that the conclusion of the theorem is valid for a wider class of problems; see, e.g., the numerical examples in Section 4.

Note that as domain decomposition methods for the linear case can be regarded as a way to construct preconditioners for Krylov methods, nonlinear domain decomposition methods, such as NLRAS, can be a way to construct nonlinear preconditioners for the Newton method applied to (2.2), and for which the following equivalent problem

$$\mathcal{F}(u) := \sum_{i=1}^P v_{i,\delta} = 0$$

is solved. However, we do not pursue this idea here, and we refer the interested reader to, e.g., [9, 10, 17, 29].

3. General asynchronous iterations. Asynchronous iterations refer to a class of parallel iterative procedures in which each processor executes its computations at the next iteration without waiting for the others to finish theirs. In other words, there is no synchronization, and the concept of iteration loses its meaning since different processors update elements of the global approximation at different times (below named timestamps). As a consequence, some processors use information which may have been updated in some processors more times than in others.

Mathematical models of asynchronous iterations were developed in order to study their convergence; see, e.g., [1, 3, 5, 22, 34, 36]. Sometimes these models are referred to as mathematical descriptions of computational models of asynchronous iterations. We review one of such models from [3], which is now classic.

Let $U = U_1 \times \dots \times U_P$, $T : U \mapsto U$, and $T_i(u) = (Tu)_i$. The goal is to solve (2.2) by means of the equivalent fixed point problem $u = T(u)$ in parallel, so that processor i runs iterations on the i th local problem, i.e., $u_i^{new} = T_i(u)$. Let us call a *timestamp* the instant of time at which at least one processor finishes its computation and updates its associated local variables. Define by t_k the sequence of all these timestamps and by $s_j^i(k)$ the sequence of integers that represent the timestamp index of the local variables coming from processor j and that are available to processor i at time instant when they start solving the local problem which is completed at time t_k . Let us also denote by I^k a subset of $\{1, \dots, P\}$ that defines the list of subdomains that are being updated at timestamp t_k .

To solve $u = T(u)$, we generate asynchronously a sequence of vectors $\{u_i^k\}_{k \geq 0}$, for $i = 1, \dots, P$, satisfying

$$(3.1) \quad u_i^{k+1} = \begin{cases} T_i(\dots, u_j^{s_j^i(k)}, \dots), & \text{for } i \in I^k, \\ u_i^k, & \text{for } i \notin I^k, \end{cases}$$

that is, either the i th portion of the solution is updated with the results of the computations in the i th processor at the timestamp t_k , or it is not updated. We also make the following three assumptions:

- (i) $\forall i, j \in \{1, \dots, P\}, \forall k \in \mathbb{N}, s_j^i(k) \leq k,$
- (ii) $\forall i, j \in \{1, \dots, P\}, \lim_{k \rightarrow \infty} s_j^i(k) = +\infty,$ and
- (iii) $\forall i \in \{1, \dots, P\}, |\{k \geq 0: i \in I^k\}| = +\infty.$

The first assumption makes sure that no future updates are being used at the beginning of the computation. The second ensures that eventually each processor is receiving new information. The last one indicates that no processor stops being updated. Convergence results for (3.1) were discussed and analyzed in several papers; see, e.g., [1, 22, 34] and references therein.

We list below three convergence results for general asynchronous iterations, which we use in our analysis of the methods discussed in this paper.

THEOREM 3.1 ([22, Theorem 4.1] following [3]). *Let T be an $n \times n$ matrix, i.e., representing a linear map. Assume that the conditions (i)–(iii) hold. If the spectral radius satisfies $\rho(|T|) < 1$, where $|T|$ has components $|t_{ij}|$, then the asynchronous iteration (3.1) converges to the solution of the fixed point problem $u = Tu$.*

For example, when T is the RAS operator (2.4) and A is a nonsingular M -matrix, we know from Theorem 2.1 that there exists a vector $w > 0$ such that $\|T_{A,\delta}\|_w < 1$. This implies that $\rho(|T_{A,\delta}|) < 1$, and we can use Theorem 3.1 to show that asynchronous RAS converges.

THEOREM 3.2 ([22, Theorem 4.4] following [19]). *Assume that u^* , the solution of (2.2), i.e., the fixed point of T , lies in the interior of U and that T is Fréchet differentiable at u^* . Assume that the conditions (i)–(iii) hold and that $\rho(|T'(u^*)|) < 1$. Then, there exists a neighborhood \mathcal{U} of u^* , such that the asynchronous iteration (3.1) converges to u^* for any $u^0 \in \mathcal{U}$.*

THEOREM 3.3 ([22, Theorem 3.3] following [19]). *Let T^k be a sequence of (nonlinear) operators, $T^k : U \rightarrow U$, such that there exists a common fixed point $u^* \in U$ for which $T^k(u^*) = u^*$ for all k . Moreover, assume that there exist $0 \leq \zeta < 1$ and $w \in \mathbb{R}^n, w > 0$, such that*

$$\|T^k(u) - u^*\|_w \leq \zeta \|u - u^*\|_w,$$

for all $k \geq 0$. Assume that the conditions (i)–(iii) hold. Then, the asynchronous iteration

$$u_i^{k+1} = \begin{cases} T_i^k(\dots, u_j^{s_j^i(k)}, \dots), & \text{for } i \in I^k, \\ u_i^k, & \text{for } i \notin I^k, \end{cases}$$

converges to u^* , the unique fixed point of all the operators T^k .

4. Asynchronous nonlinear RAS (ANLRAS). Asynchronous parallel methods for nonlinear equations have been studied before; see, e.g., [2, 6, 15, 20, 25, 32, 35, 37]. The experimental results obtained in these papers showed that the asynchronous methods perform better in terms of execution times than their synchronous counterparts; see, e.g., [15, Tables C5–C7]. Here we present the asynchronous RAS method for nonlinear problems, provide local convergence results, and compare its performance with that of the synchronous counterpart. We emphasize that our approach is well suited to heterogeneous networks and in general to parallel computers where synchronization and communication are at a premium.

The asynchronous NLRAS method consists of the standard asynchronous iteration (3.1), where T_i is given by (2.5) and (2.7); see also Algorithm 1 below. We begin with a local convergence theorem.

THEOREM 4.1. *Let u^* be the solution of (2.2). Suppose that F' exists in a neighborhood of u^* and that $F'(u^*)$ is a nonsingular M -matrix. Assume that the conditions (i)–(iii) hold. Then, there exists a neighborhood \mathcal{U} of u^* such that the asynchronous NLRAS defined by (3.1) with T_i given by (2.5) and (2.7), converges to u^* for any initial vector $u^0 \in \mathcal{U}$.*

Proof. The proof is essentially the same as the one of Theorem 2.2. However, here we need the additional hypothesis that T is differentiable in a region containing u^* . We recall that the NLRAS iteration is given by

$$u^{k+1} = T(u^k), \quad T(u) = u + \sum_{i=1}^P R_{i,0}^\top v_{i,\delta},$$

where $F_{i,\delta}(u + R_{i,\delta}^\top v_{i,\delta}) = 0$. To prove that the asynchronous iteration associated with the operator T converges, it suffices to prove that $\|T'(u^*)\| < 1$ for some operator norm. To this end, we compute

$$T'(u) = I + \sum_{i=1}^P R_{i,\delta}^\top \frac{\partial v_{i,\delta}}{\partial u}.$$

Also, differentiating $F_{i,\delta}$, we obtain

$$F'_{i,\delta}(u + R_{i,\delta}^\top v_{i,\delta}) \left(I + R_{i,\delta}^\top \frac{\partial v_{i,\delta}}{\partial u} \right) = 0,$$

hence,

$$(4.1) \quad \begin{aligned} F'_{i,\delta}(u + R_{i,\delta}^\top v_{i,\delta}) R_{i,\delta}^\top \frac{\partial v_{i,\delta}}{\partial u} &= -F'_{i,\delta}(u + R_{i,\delta}^\top v_{i,\delta}), \\ \frac{\partial v_{i,\delta}}{\partial u} &= - (R_{i,\delta} F'(u + R_{i,\delta}^\top v_{i,\delta}) R_{i,\delta})^{-1} F'_{i,\delta}(u + R_{i,\delta}^\top v_{i,\delta}). \end{aligned}$$

Thus,

$$(4.2) \quad T'(u) = I - \sum_{i=1}^P R_{i,\delta}^\top (R_{i,\delta} F'(u + R_{i,\delta}^\top v_{i,\delta}) R_{i,\delta})^{-1} F'_{i,\delta}(u + R_{i,\delta}^\top v_{i,\delta}).$$

Evaluating (4.2) at $u = u^*$, we obtain

$$T'(u^*) = I - \sum_{i=1}^P R_{i,\delta}^\top (R_{i,\delta} F'(u^*) R_{i,\delta})^{-1} F'_{i,\delta}(u^*) = T_{F'(u^*),\delta},$$

where the latter is exactly the same as in (2.4) but applied to $F'(u^*)$. Hence, it holds that $\|T'(u^*)\|_w = \|T_{F'(u^*),\delta}\|_w < 1$ for $w > 0$ such that $F'(u^*)w > 0$. Using Theorem 3.2 with the observation that $\rho(|T_{F'(u^*),\delta}|) \leq \|T_{F'(u^*),\delta}\|_w$, we conclude that there exists a neighborhood of \mathcal{U} of u^* such that the asynchronous NLRAS is convergent. \square

As was the case for the synchronous method, the convergence proof relies on the hypothesis that $F'(u^*)$ is a nonsingular M -matrix. Many of the problems in practice do not satisfy this hypothesis, but as we illustrate with the numerical experiments below, the conclusion of the theorem still holds in practice.

We describe next the asynchronous stopping criterion we used in our implementation of asynchronous NLRAS on a shared memory machine. For the sake of simplicity, we assume that each computational thread is assigned to a unique subdomain. The goal is to reduce the relative residual norm below a prescribed tolerance $\epsilon > 0$, namely

$$(4.3) \quad \|F(u^k)\|_2 / \|F(u^0)\|_2 < \epsilon.$$

We observe that the L^2 -norm of the global residual can be computed as a sum of parts

$$\|F(u)\|_2^2 = \sum_{i=1}^P \|F_{i,0}(u)\|_2^2.$$

Hence, in order to verify that the global stopping criterion was achieved, it suffices to verify that the local solution on each subdomain satisfies the local convergence criteria $\|F_{i,0}(u)\|_2/\|F(u^0)\|_2 < \epsilon/\sqrt{P}$. This condition ensures that the contribution of the local residual is such that (4.3) holds. In our implementation, a global boolean variable is declared and is then set to true when all the subdomains satisfy the local convergence criteria. We comment that in order to update the local iteration vectors and residuals, we only need to communicate with the neighboring subdomains. This allows us to test for the local stopping criteria and hence avoid global communication between the subdomains/processors.

We describe the implementation of the asynchronous NLRAS in Algorithm 1.

Algorithm 1 Asynchronous nonlinear RAS (ANLRAS)

Require: Initial vector u^0 , tolerance ϵ

- 1: Compute $\|F(u^0)\|_2$
- 2: Set `global_convergence=false`, and `local_convergence[k]=false`, $k = 1, \dots, P$
- In parallel, each processor i:**
- 3: **while not** `global_convergence` **do**
- 4: Compute $v_{i,\delta}$ by solving (2.5) using Newton's method
- 5: Update $u = u + R_{i,0}^\top v_{i,\delta}$
- 6: **if** $\|F_{i,0}(u)\|_2/\|F(u^0)\|_2 < \epsilon/\sqrt{P}$ **then**
- 7: Set `local_convergence[i]=true`
- 8: **end if**
- 9: **if** $i==1$ **then**
- 10: **if** `local_convergence[k]`, $\forall k = 1, \dots, P$ **then**
- 11: Set `global_convergence=true`
- 12: **end if**
- 13: **end if**
- 14: **end while (for processor i)**
- 15: **Output:** u^*

We mention that once a processor/subdomain achieves local convergence, the local convergence can be lost before the global convergence flag is triggered. This is why we compute the global residual norm again after the iterations stopped and verify that it is below the convergence threshold. This was always satisfied in our experiments.

We illustrate the performance of this algorithm on two nonlinear PDEs. Let us first consider the following problem:

$$(4.4) \quad -\nabla^2 u + g(u) = f \quad \text{on } \Omega, \quad u = 0 \quad \text{on } \partial\Omega,$$

where $\Omega = (0, 1)^2$ and $g(u) = ue^u$. We choose f such that $\sin(\pi x)\sin(\pi y)$ is the exact solution of (4.4). The discretization of (4.4) using standard 5-points stencil with mesh size h on the nodes (ih, jh) yields the following nonlinear system of equations

$$(4.5) \quad (F(u))_{i,j} = \frac{1}{h^2}(4u_{i,j} - u_{i-1,j} - u_{i+1,j} - u_{i,j+1} - u_{i,j-1}) + u_{i,j}e^{u_{i,j}} - f_{i,j} = 0,$$

with $u_{0,j} = u_{N+1,j} = u_{j,0} = u_{j,N+1} = 0$. We present and analyze computational results for parallel asynchronous and synchronous NLRAS for the discrete nonlinear system of

equations (4.5). The computational experiments were carried out on a shared memory machine with 88 CPU cores/176 threads and 1536GB of RAM. The parallelization of the solvers was implemented in C++ using the OpenMP multithreading library. The linear algebra data structures and solvers were provided by EIGEN [30]. The local problems in (2.5) were solved using an inner Newton’s method where the local convergence criterion is when the norm of the difference between two updates in the local Newton iterations is smaller than 10^{-10} , and the local Jacobian matrices arising from the inner Newton iterations were computed analytically. We use a checker-board partitioning for the domain $\Omega = (0, 1)^2$ with a total of 10 000 discretization points. We choose the zero vector as an initial vector for both the RAS iteration and the local Newton iterations.




FIG. 4.1. Convergence curves of synchronous and asynchronous NLRAS for problem (4.4) with $P = 16$. Left: homogeneous network. Right: one processor slower than the others by a factor of two.

For each tolerance in $\{10^{-1}, 10^{-2}, \dots, 10^{-8}\}$, we run Algorithm 1 and measure the corresponding computational time (in seconds) required to reach that specific tolerance. We then plot the timing versus the relative residual norm for both methods as shown in Figure 4.1 (left). We observe in this case that the synchronous method is faster in terms of execution time. We repeat the same experiment, but with one processor twice as slow as the others by performing the inner Newton solve twice for one processor, and we report the result in Figure 4.1 (right). It is clear that in this case, the asynchronous method outperforms the synchronous one in terms of execution time. This indicates that even with a small number of processors, if the network is heterogeneous, asynchronous methods may be preferred. The same effect would be observed if a large load imbalance is present, that is, if one or more processors need to perform more work than others.

We report in Table 4.1 the number of local updates, varying the number of subdomains for both the synchronous and the asynchronous NLRAS. With one processor twice as slow, the asynchronous method is faster than the synchronous method even if, on average, the asynchronous method requires more local updates.

We also compare the number of Newton solves required for each subdomain/processor for both the synchronous and asynchronous cases. We illustrate this in Figure 4.2, where we see that the number of Newton iterations is slightly different for each processor. We also observe that the asynchronous case requires more Newton iterations compared to the synchronous case since it usually requires more local updates.

TABLE 4.1

Number of iterations/updates and computational time (in seconds) for problem (4.4) where one processor is twice as slow.

	#DOF	#subdomains	NLRAS	
			#iter	time
			(min,mean,max)	
Synchronous				
	2 500	4	160	2.40
	5 625	9	323	4.79
	10 000	16	554	8.55
Asynchronous				
	2 500	4	(149, 250, 287)	2.18
	5 625	9	(282, 508, 570)	4.20
	10 000	16	(470, 834, 892)	6.89




FIG. 4.2. Total number of Newton iterations per processor for problem (4.4) with $P = 16$.

In order to investigate the influence of the heterogeneity factor on the performance of both the synchronous and asynchronous methods, we perform an experiment where we vary the heterogeneity factor in the interval $[0, 4]$ and compute the timing required to reach a tolerance of 10^{-8} . This can be done by enforcing one processor to wait a given amount of time using the C++ built-in function `sleep_for()` that is proportional to the local Newton solve. The plots in Figure 4.3 illustrate the time required to reach convergence for each method and with different heterogeneity factors. We can observe that there is a heterogeneity threshold for which the asynchronous method starts to outperform the synchronous one. We can also see that the time required for the asynchronous method remains constant while the time required for the synchronous method increases roughly proportionally with the heterogeneity factor.

For our next set of experiments, we consider a problem where the nonlinearity is on the boundary. This is a temperature control model defined as follows:

$$\begin{aligned}
 &-\nabla^2 u + v \cdot \nabla u = f \quad \text{on } \Omega \\
 (4.6) \quad &\frac{\partial u}{\partial \nu} + \varphi(u) = 0 \quad \text{on } \Gamma \quad \text{and} \quad u = 0 \quad \text{on } \partial\Omega \setminus \Gamma,
 \end{aligned}$$

where $v \in \mathbb{R}^2$, $\Gamma \subset \partial\Omega$, and $\varphi: \mathbb{R} \mapsto \mathbb{R}$ is a function of the form $\varphi(u) = \gamma \log(\beta + \alpha u)$, with $\alpha, \beta, \gamma > 0$. For our experiments we consider $\Omega = (0, 1)^2$, $\Gamma = \{0\} \times [0, 1]$, and $v^T = [1, 1]$. We discretize (4.6) using finite differences with $n = 10\,000$ discretization points.




FIG. 4.3. Execution time of synchronous and asynchronous NLRAS with different heterogeneity factors for problem (4.4) with $P = 16$.

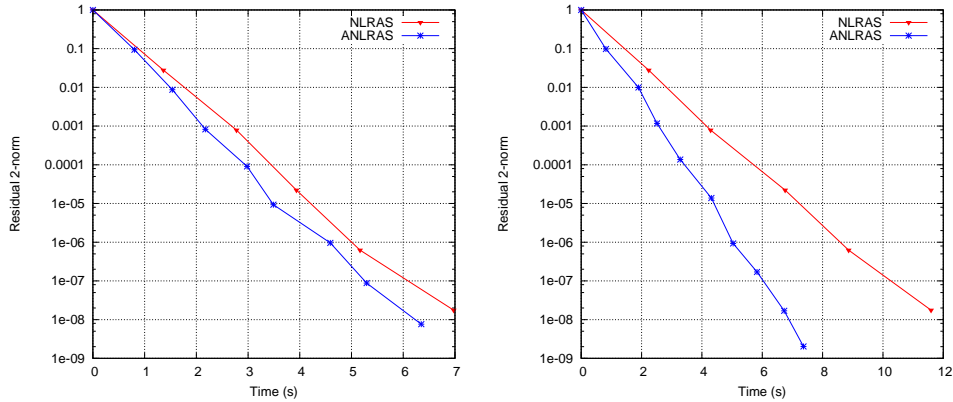


FIG. 4.4. Convergence curves of synchronous and asynchronous NLRAS for problem (4.6) with $\alpha = 1$, $\beta = 1$, $\gamma = 10^2$. Left: homogeneous processor speed. Right: one processor is twice as slow.

The resulting convergence plots are presented in Figure 4.4, and they illustrate the convergence of synchronous and asynchronous NLRAS for (4.6) with $P = 16$ subdomains. We can clearly see that in this case the asynchronous method is more advantageous, even in the case of an homogeneous network. This is due to the load imbalance among subdomains solves. Some subdomains (those lying in the boundary of Ω) require a nonlinear solver, while others require only a linear solver. The advantage is enhanced when one of the processors is twice as slow as the others.

5. Two-level nonlinear RAS (2L-NLRAS). In order to improve the convergence of NLRAS and ensure the scalability of the method, an additional coarse correction can be incorporated. Since the problem is nonlinear, we need to use a nonlinear coarse grid correction. We opt here for the full approximation scheme (FAS) method to perform the coarse grid correction; see, e.g., [8]. Let V_0 denote the coarse space and $R_0 : V \mapsto V_0$ the coarse restriction operator. We can then define by means of R_0 the coarse nonlinear problem $F_0 : V_0 \mapsto V_0$ which could be stated by using a coarse discretization of (2.1) or using a

Galerkin approach, namely

$$F_0(u) = R_0 F(R_0^\top u).$$

The coarse correction $v_0 \in V_0$ using the FAS is defined as the solution of the nonlinear problem

$$(5.1) \quad F_0(v_0 + R_0 u) = F_0(R_0 u) - R_0 F(u),$$

which can be incorporated into (2.6) either in an additive or multiplicative manner. Following [26], here we opt to use the additive variant. In order to prevent overcorrection on the overlap, we use a weighted two-level method of the form

$$(5.2) \quad u^{k+1} = u^k + \frac{1}{2} \sum_{i=1}^P R_{i,0}^\top v_{i,\delta}^k + \frac{1}{2} R_0^\top v_0^k;$$

see [24] for an alternative approach. FAS [7] is a particular case of the nonlinear multilevel method (NMLM) of Hackbusch [27]. Although NMLM is a globally convergent iteration [27], there exist no such result for FAS. However, there exist cases for which the convergence of FAS was proven for a class of mildly nonlinear PDEs [33, Theorem 5.1].

In the rest of the section, we discuss the local convergence of the two-level NLRAS described in (5.2) and present numerical experiments which illustrate the benefits of the two-level approach.

Let us thus define the mapping \tilde{T} corresponding to the fixed point iteration (5.2), i.e.,

$$(5.3) \quad \tilde{T}(u) = u + \frac{1}{2} \sum_{i=1}^P R_{i,0}^\top v_{i,\delta} + \frac{1}{2} R_0^\top v_0.$$

In order to show that (5.2) converges to the solution of (2.2), it suffices to show that there exists a norm such that $\|\tilde{T}'(u^*)\| < 1$. We have that

$$(5.4) \quad \tilde{T}'(u) = I + \frac{1}{2} \sum_{i=1}^P R_{i,0}^\top \frac{\partial v_{i,\delta}}{\partial u} + \frac{1}{2} R_0^\top \frac{\partial v_0}{\partial u}.$$

Moreover, we know that (4.1) holds. It remains to compute the derivative corresponding to the coarse correction v_0 . Indeed, taking the derivative on both sides of (5.1) with respect to u , we obtain

$$\begin{aligned} F_0'(v_0 + R_0 u) \left(\frac{\partial v_0}{\partial u} + R_0 \right) &= F_0'(R_0 u) R_0 - R_0 F'(u) \\ \frac{\partial v_0}{\partial u} + R_0 &= (F_0'(v_0 + R_0 u))^{-1} (F_0'(R_0 u) R_0 - R_0 F'(u)) \\ \frac{\partial v_0}{\partial u} &= (F_0'(v_0 + R_0 u))^{-1} (F_0'(R_0 u) - I) R_0 \\ &\quad - F_0'(v_0 + R_0 u)^{-1} R_0 F'(u). \end{aligned}$$

Substituting both terms $\frac{\partial v_{i,\delta}}{\partial u}$ and $\frac{\partial v_0}{\partial u}$ in (5.4), we obtain that the derivative of \tilde{T} at $u = u^*$ reduces to

$$(5.5) \quad \begin{aligned} I - \frac{1}{2} \sum_{i=1}^P R_{i,0}^\top (R_{i,\delta} F'(u^*) R_{i,\delta})^{-1} R_{i,\delta} F'(u^*) \\ - \frac{1}{2} R_0^\top (R_0 F'(R_0^\top R_0 u^*) R_0^\top)^{-1} R_0 F'(u^*). \end{aligned}$$

In order to prove convergence of this two-level method, it suffices then to show that $\tilde{T}'(u^*)$ as in (5.5) is a contraction. We can use the same techniques as for the one-level method in Theorem 2.2. One added complication is that unlike, e.g., in [26], for our coarse grid restriction $R_0^\top R_0 \neq I$, and thus additional hypotheses are needed. Then, we can use arguments similar to those in the proof of [4, Theorem 7.3] to complete the proof, as we show below. Let

$$B = \sum_{i=1}^P R_{i,0}^\top (R_{i,\delta} F'(u^*) R_{i,\delta}^\top)^{-1} R_{i,\delta}$$

and

$$B_0 = R_0^\top (R_0 F'(R_0^\top R_0 u^*) R_0^\top)^{-1} R_0.$$

THEOREM 5.1. *Let u^* be the solution of (2.2). Assume that F' exists in a neighborhood of u^* and that $F'(u^*)$ and $F'(R_0^\top R_0 u^*)$ are nonsingular M -matrices. Assume further that there exist a weak regular splitting $F'(u^*) = M_0 - N_0$, and a diagonal matrix E_0 such that*

$$0 \leq E_0 \leq I,$$

and

$$B_0 = E_0 M_0^{-1}.$$

Let $w > 0$ be such that $F'(u^*)w > 0$. Then $\rho(\tilde{T}'(u^*)) < \|\tilde{T}'(u^*)\|_w < 1$, and the iteration (5.2) converges linearly to u^* .

Proof. We start by showing that $\tilde{T}' \geq 0$. Indeed, we observe that

$$\tilde{T}'(u^*) = \frac{1}{2} \underbrace{(I - BF'(u^*))}_{\geq 0} + \frac{1}{2} \underbrace{E_0(I - M_0^{-1}F'(u^*))}_{\geq 0} + \frac{1}{2} \underbrace{(I - E_0)}_{\geq 0}.$$

Now, let us re-write the local contraction operator at the root $\tilde{T}'(u^*)$ (5.5) as

$$\tilde{T}'(u^*) = I - \frac{1}{2}(B + B_0)F'(u^*).$$

We remark that as B and B_0 are restrictions and prolongations of nonnegative matrices, it follows that wherever the restriction and prolongation operators are nonnegative, these matrices are nonnegative. Furthermore, $B \geq 0$ is nonsingular, in fact, it is a RAS preconditioner. Therefore, it cannot have a zero row, and thus $Be > 0$, for any positive vector e . Let $e = F'(u^*)w > 0$. Then, we have that $Be > 0$ and $B_0e \geq 0$, so that

$$\tilde{T}'(u^*)w = w - \frac{1}{2}(Be + B_0e) < w,$$

concluding that $\|\tilde{T}'(u^*)\|_w < 1$. \square

We note that the hypotheses mentioned in Theorem (5.1) are not always satisfied. In particular, for our coarse grid, the nonnegativity of the operator \tilde{T}' does not necessarily hold. For the particular case where the coarse grid corresponds to a boolean matrix, it is shown in [21] that the choice of $E_0 = R_0^\top R_0$ is possible.

We present now several numerical experiments comparing the performance of two-level NLRAS with the one-level counterpart for problem (4.4), and conduct scaling experiments by varying the number of subdomains.

We show in Figure 5.1 a weak scaling experiment for the problem (4.4). We compare the one- and two-level NLRAS for a 2×2 and 3×3 decomposition. The total number of degrees of freedom (DOFs) in each subdomain is kept the same for both cases, namely 625 variables per subdomain. As we can observe, the convergence of the one-level NLRAS deteriorates when the number of subdomains increases from $P = 4$ to $P = 9$. In contrast, the convergence rate of the two-level NLRAS remains approximately the same. We summarize weak scalability




FIG. 5.1. Weak scaling experiment for one- and two-level NLRAS for problem (4.4) using a 2×2 checker-board partition (left) and a 3×3 checker-board partition (right).

results with increasing numbers of subdomains for iteration (5.2) in Table 5.1, where we report the number of iterations to reach a tolerance of 10^{-8} . These results show that the method without coarse correction fails to converge in 100 iterations, while for the two-level NLRAS, the number of iterations has a slow growth with the number of subdomains.

TABLE 5.1

Number of iterations performed by NLRAS and two-level NLRAS for the problem (4.4) for different numbers of subdomains P and a tolerance of $\epsilon = 10^{-8}$.

P	$ V $	$ V_0 $	NLRAS	2L-NLRAS
4	2 500	36	160	62
9	5 625	81	323	64
16	10 000	144	554	94
25	15 625	225	853	85
36	22 500	324	1221	82
49	30 625	441	1657	80

We repeat the same weak-scaling experiment but for a nonlinear diffusion equation problem of the form

$$(5.6) \quad \begin{aligned}
 -\nabla \cdot ((1 + u^2)\nabla u) &= f(x, y), & \text{in } \Omega = (0, 1)^2, \\
 u(x, y) &= 0, & \text{on } \partial\Omega.
 \end{aligned}$$

We summarize in Table 5.2 the convergence results of iteration (5.2) for this problem. We point out that for a discretization of (5.6) using finite differences, $F'(u)$ is a nonsingular M -matrix for $u \in \Omega$, and therefore the hypothesis of our local convergence theorems are satisfied in this

case. The results in Table 5.2 are similar to those in Table 5.1. The addition of the coarse correction makes the algorithm scalable in a sense that the number of iterations required to achieve the prescribed tolerance has a very slow growth as the number of subdomains increases. Comparing Tables 5.1 and 5.2, we note that for two problems with different type of nonlinearity, the two-level method behaves similarly in terms of number of iterations to convergence.

TABLE 5.2
 Number of iterations performed by NLRAS and two-level NLRAS for problem (5.6) for different number of subdomains P and a tolerance of $\epsilon = 10^{-8}$.

P	$ V $	$ V_0 $	NLRAS	2L-NLRAS
4	2 500	36	183	63
9	5 625	81	373	64
16	10 000	144	642	97
25	15 625	225	992	89
36	22 500	324	1411	84
49	30 625	441	1913	85

We end this section with two practical observations on the implementation of this two-level NLRAS method. The coarse correction step in (5.1) requires the use of Newton’s iteration, which in turn necessitates the knowledge of the action of $F(u)$ as well as $F'(u)$. These quantities can be obtained by means of the local functions $F_{i,\delta}$ as follows. By recalling the definition of $F_{i,\delta}$, we have that

$$F_0(u) = R_0 F(R_0^\top u) = R_0 \left(\sum_{i=1}^P R_{i,0}^\top F_{i,\delta}(R_0^\top u) \right).$$

Similarly, we obtain for the derivative of F_0

$$F'_0(u) = R_0 \left(\sum_{i=1}^P R_{i,0}^\top F'_{i,\delta}(R_0^\top u) \right) R_0^\top.$$

We comment on the appropriate choice of the coarse grid in this case to ensure convergence and scalability of the two-level NLRAS. The one-level NLRAS has a residual $F(u)$, which is zero outside the overlap. Indeed, the global nonlinear residual is defined as $F(u^k)$, where u^k is the current approximation of the solution of (2.2). Moreover,

$$R_{i,\delta} F(u^k) = R_{i,\delta} F(u^{k-1} + \sum_{j=1}^P R_{j,0}^\top v_{j,\delta}^{k-1}) = F_{i,\delta}(u^{k-1} + R_{i,\delta}^\top v_{i,0}^{k-1} + \sum_{j \neq i} R_{j,0}^\top v_{j,\delta}^{k-1}).$$

Using the fact that $v_{i,\delta}^{k-1}$ satisfies $F_{i,\delta}(u^{k-1} + R_{i,0}^\top v_{i,\delta}^{k-1}) = 0$, it follows that the local nonlinear residual $R_{i,\delta} F(u^k)$ is zero outside the overlap. This shows that a good coarse grid should use this information for the construction of the coarse space. Thus, the coarse grid should be chosen on the overlap.

6. Asynchronous two-level nonlinear RAS (A2L-NLRAS). In this section, we study the asynchronous version of the iteration (5.2). We present related implementation details, local convergence theory, and numerical experiments. The information for the coarse grid correction comes from all subdomains, but it cannot be used until the information from all

subdomains has arrived. Only then, the correction is added, and the process repeats. To implement this, we follow some ideas described in [26] for linear problems. More in detail, we use a weighted additive NLRAS method (see [24] for multiplicative variants), and in order to prevent the coarse correction from over-correcting, we add the coarse grid only when all the subdomains have sent new data. This can be realized by defining a local boolean variable for each thread/subdomain. The coarse grid is then added when all these local boolean variables are set to true, and this is done without a synchronization point. We describe the asynchronous two-level NLRAS in Algorithm 2.

Algorithm 2 Asynchronous two-level nonlinear RAS (A2L-NLRAS)

Require: Initial vector u^0 , tolerance ϵ

- 1: Compute $\|F(u^0)\|_2$
- 2: Set `global_convergence=false`, and `local_convergence[k]=false`, $k = 1, \dots, P$
- 3: Set `local_update[k]=false`, and `local_correction[k]=false`, $k = 1, \dots, P$
- In parallel, each processor i :**
- 4: **while not** `global_convergence` **do**
- 5: **if** $i > 0$ **then**
- 6: Compute $v_{i,\delta}$ by solving (2.5) using Newton's method
- 7: **if** `local_correction[i]` **then**
- 8: Update $u = u + \frac{1}{2}R_{i,0}^\top v_{i,\delta} + \frac{1}{2}R_{i,0}^\top R_{i,\delta} v_0$
- 9: Set `local_correction[i]=false`
- 10: **else**
- 11: Update $u = u + R_{i,0}^\top v_{i,\delta}$
- 12: Set `local_update[i]=true`
- 13: **end if**
- 14: **if** $\|F_{i,0}(u)\|_2 / \|F(u^0)\|_2 < \epsilon / \sqrt{P}$ **then**
- 15: Set `local_convergence[i]=true`
- 16: **end if**
- 17: **if** $i==1$ **then**
- 18: **if** `local_convergence[k], $\forall k = 1, \dots, P$` **then**
- 19: Set `global_convergence=true`
- 20: **end if**
- 21: **end if**
- 22: **else**
- 23: **if** `local_update[k], $\forall k = 1, \dots, P$` **then**
- 24: Compute the coarse correction v_0 using (5.1)
- 25: Set `local_correction[k]=true`, $k = 1, \dots, P$
- 26: Set `local_update[k]=false`, $k = 1, \dots, P$
- 27: **end if**
- 28: **end if**
- 29: **end while (for processor i)**
- 30: **Output:** u^*

We state and prove a result on the local convergence of two-level NLRAS as presented in Algorithm 2.

THEOREM 6.1. *Let u^* be the solution of (2.2). Assume that F' exists in a neighborhood of u^* and that $F'(u^*)$ and $F'(R_0^\top R_0 u^*)$ are nonsingular M -matrices. Let $w > 0$ be such that $F'(u^*)w > 0$. Assume that the conditions (i)–(iii) hold. Then Algorithm 2 converges to u^* , the solution of (2.2).*

Proof. We will use Theorem 3.3. To that end, we define a sequence of operators T^k , which is either

$$(6.1) \quad T^k(u) = \tilde{T}(u) = u + \frac{1}{2} \sum_{i=1}^P R_{i,0}^\top v_{i,\delta} + \frac{1}{2} R_0^\top v_0,$$

if the coarse grid in Algorithm 2 is added, or

$$(6.2) \quad T^k(u) = T(u) = u + \sum_{i=1}^P R_{i,0}^\top v_{i,\delta},$$

otherwise. Of course the operators (6.1) are those of the two-level NLRAS method (5.3), while the ones in (6.2) are those of the one-level method (2.7). It follows that both operators have u^* as a common fixed point, and in both cases $\|T^k\|_w < 1$, by Theorems 2.2 and 5.1. Thus, from Theorem 3.3, Algorithm 2 converges to u^* . \square




FIG. 6.1. Convergence curves of synchronous and asynchronous two-level NLRAS for problem (4.4) with $P = 16$. Left: homogeneous network. Right: one processor four times slower than the others.

We compare numerically the parallel synchronous and asynchronous two-level NLRAS for the nonlinear problem (4.4). We keep the same setting used in Section 4. In particular, the Jacobians are computed exactly for the inner Newton solves. The coarse grid is chosen using a full weighting restriction matrix. We use 10 000 discretization points for the domain $\Omega = (0, 1)^2$ which we partition into 16 subdomains. The total size of the coarse grid is 144 discretization points. We test our methods in two settings: first with all processors of the same speed, and then with one processor slower than the others, in this case, four times slower. These experiments are reported in Figure 6.1. We can clearly see that in the case of an inhomogeneous network, the asynchronous two-level method is faster than its synchronous counterpart. For our weak scaling study, varying the number of subdomains, we report in Figure 6.2 the time required to reach a tolerance of $\epsilon = 10^{-8}$ for the synchronous and asynchronous two-level NLRAS for problems (4.4) and (5.6). We keep the computation unbalanced, where one processor is four times slower than the others. Once again, we see that the asynchronous method converges faster than the synchronous one.

We conclude with a table comparing execution times and the number of processor updates for the one-level and two-level methods in both the synchronous and asynchronous versions




FIG. 6.2. Convergence time in seconds of synchronous and asynchronous two-level NLRAS versus the number of subdomains for $\epsilon = 10^{-8}$ where one processor is four times slower than the others. Left: problem (4.4). Right: problem (5.6).

for problem (4.4) in 2D and 3D using $P + 1$ processors. We note that each subdomain is assigned to a processor and the coarse grid is computed on a separate one. In the 3D case, the domain $\Omega = (0, 1)^3$ is decomposed into P equally sized bricks, and their number is the same in each direction. The size of each local problem is 1000×1000 . We assume that one processor is four times slower than the others. The convergence results are reported in Table 6.1. We see from Table 6.1 that the asynchronous method outperforms the synchronous NLRAS. The addition of the coarse grid correction makes both methods even faster, with an advantage for the asynchronous one. The results in Table 6.1 also show that the timing required for the coarse step increases with the number of subdomains.

Observe that, as expected, the asynchronous method has a larger number of updates on average than the number of iterations of the synchronous method. Nevertheless, its execution time is faster, confirming that it can outperform the synchronous method. This is true for both the one- and two-level methods.

TABLE 6.1

Time (in sec) and the number of iterations or average number of updates for synchronous and asynchronous NLRAS and its two-level variant for problem (4.4) in 2D (top) and 3D (bottom) to reach a tolerance of $\epsilon = 10^{-8}$ for different numbers of subdomains with constant local problem size.

P	V	V ₀	NLRAS				2L-NLRAS					
			sync		async		sync		async			
			time	#iter	time	#updates	time	coarse	#iter	time	coarse	#updates
4	2 500	36	4.44	160	3.38	374	1.71	0.20	62	1.39	0.57	131
9	6 525	81	9.67	323	5.22	581	1.87	0.45	64	1.40	0.72	153
16	10 000	144	16.42	554	7.18	880	2.86	1.30	94	1.47	0.96	170
25	15 625	225	26.15	853	10.98	1339	2.82	1.66	85	1.68	1.25	187
8	8 000	125	21.53	72	17.28	163	17.64	1.68	56	14.17	4.02	150
27	27 000	343	47.51	147	30.71	268	24.23	9.55	77	17.71	5.82	159
64	64 000	1000	56.73	251	41.99	447	17.86	17.71	67	16.70	13.51	182

7. Conclusion. In summary, the new parallel nonlinear RAS algorithm converges to the unique solution of the problem in the original domain of the PDE. An asynchronous version can be about twice as fast or faster, depending on the inhomogeneities of the computation and the communication costs. The two-level version can be an order of magnitude faster, and its asynchronous version is faster, still exhibiting slow growth in execution times when the number of processors increases in weak scaling.

REFERENCES

- [1] J. M. BAHİ, S. CONTASSOT-VIVIER, AND R. COUTURIER, *Parallel Iterative Algorithms*, Chapman & Hall, Boca Raton, 2008.
- [2] Z.-Z. BAI, V. MIGALLÓN, J. PENADÉS, AND D. B. SZYLD, *Block and asynchronous two-stage methods for mildly nonlinear systems*, *Numer. Math.*, 82 (1999), pp. 1–20.
- [3] G. M. BAUDET, *Asynchronous iterative methods for multiprocessors*, *J. Assoc. Comput. Mach.*, 25 (1978), pp. 226–244.
- [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] D. P. BERTSEKAS AND J. N. TSITSIKLIS, *Parallel and Distributed Computation: Numerical Methods*, Prentice-Hall, Englewood Cliffs, 1989.
- [6] A. BHAYA, E. KASZKUREWICZ, AND F. MOTA, *Asynchronous block-iterative methods for almost linear equations*, *Linear Algebra Appl.*, 154/156 (1991), pp. 487–508.
- [7] A. BRANDT, *Multi-level adaptive solutions to boundary-value problems*, *Math. Comp.*, 31 (1977), pp. 333–390.
- [8] ———, *Guide to multigrid development*, in *Multigrid Methods* (Cologne, 1981), W. Hackbusch and U. Trottenberg, eds., vol. 960 of *Lecture Notes in Math.*, Springer, Berlin, 1982, pp. 220–312.
- [9] X.-C. CAI AND D. E. KEYES, *Nonlinearly preconditioned inexact Newton algorithms*, *SIAM J. Sci. Comput.*, 24 (2002), pp. 183–200.
- [10] X.-C. CAI AND X. LI, *Inexact Newton methods with restricted additive Schwarz based nonlinear elimination for problems with high local nonlinearity*, *SIAM J. Sci. Comput.*, 33 (2011), pp. 746–762.
- [11] 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.
- [12] F. CHAOUQUI, M. J. GANDER, P. M. KUMBHAR, AND T. VANZAN, *Linear and nonlinear substructured Restricted Additive Schwarz iterations and preconditioning*, *Numer. Algorithms*, 91 (2022), pp. 81–107.
- [13] G. CIARAMELLA AND T. VANZAN, *Substructured two-grid and multi-grid domain decomposition methods*, *Numer. Algorithms*, 91 (2022), pp. 413–448.
- [14] L. B. ĆIRIĆ, *Generalized contractions and fixed-point theorems*, *Publ. Inst. Math. (Beograd) (N.S.)*, 12(26) (1971), pp. 19–26.
- [15] M. CHAU, D. EL BAZ, R. GUIVARCH, AND P. SPITERI, *MPI implementation of parallel subdomain methods for linear and nonlinear convection–diffusion problems*, *J. Parallel Distrib. Comput.*, 67 (2007), pp. 581–591.
- [16] M. DRYJA AND W. HACKBUSCH, *On the nonlinear domain decomposition method*, *BIT*, 37 (1997), pp. 296–311.
- [17] V. DOLEAN, M. J. GANDER, W. KHERIJI, F. KWOK, AND R. MASSON, *Nonlinear preconditioning: how to use a nonlinear Schwarz method to precondition Newton’s method*, *SIAM J. Sci. Comput.*, 38 (2016), pp. 3357–A3380.
- [18] E. EFSTATHIOU AND M. J. GANDER, *Why restricted additive Schwarz converges faster than additive Schwarz*, *BIT*, 43 (2003), pp. 945–959.
- [19] M. N. EL TARAIZI, *Some convergence results for asynchronous algorithms*, *Numer. Math.*, 39 (1982), pp. 325–340.
- [20] A. FROMMER AND H. SCHWANDT, *Asynchronous parallel methods for enclosing solutions of nonlinear equations*, *J. Comput. Appl. Math.*, 60 (1995), pp. 47–62.
- [21] A. FROMMER AND D. B. SZYLD, *Weighted max norms, splittings, and overlapping additive Schwarz iterations*, *Numer. Math.*, 83 (1999), pp. 259–278.
- [22] ———, *On asynchronous iterations*, *J. Comput. Appl. Math.*, 123 (2000), pp. 201–216.
- [23] ———, *An algebraic convergence theory for restricted additive Schwarz methods using weighted max norms*, *SIAM J. Numer. Anal.*, 39 (2001), pp. 463–479.
- [24] G. GBIKPI-BENISSAN AND F. MAGOULÈS, *Asynchronous multiplicative coarse-space correction*, *SIAM J. Sci. Comput.*, 44 (2022), pp. C237–C259.
- [25] L. GIRAUD AND P. SPITERI, *Résolution parallèle de problèmes aux limites non linéaires*, *RAIRO Modél. Math. Anal. Numér.*, 25 (1991), pp. 579–606.

- [26] C. GLUSA, E. G. BOMAN, E. CHOW, S. RAJAMANICKAM, AND D. B. SZYLD, *Scalable asynchronous domain decomposition solvers*, SIAM J. Sci. Comput., 42 (2020), pp. C384–C409.
- [27] W. HACKBUSCH, *Multigrid Methods and Applications*, Springer, Berlin, 1985.
- [28] A. HEINLEIN AND M. LANSER, *Additive and hybrid nonlinear two-level Schwarz methods and energy minimizing coarse spaces for unstructured grids*, SIAM J. Sci. Comput., 42 (2020), pp. A2461–A2488.
- [29] F.-N. HWANG AND X.-C. CAI, *Improving robustness and parallel scalability of Newton method through nonlinear preconditioning*, in Domain Decomposition Methods in Science and Engineering, R. Kornhuber, R. Hoppe, J. Périaux, O. Pironneau, O. Widlund, and J. Xu, eds., vol. 40 of Lect. Notes Comput. Sci. Eng., Springer, Berlin, 2005, pp. 201–208.
- [30] B. JACOB, GAËL GUENNEBAUD ET AL., *Eigen*, C++ template library, 2010.
<http://eigen.tuxfamily.org>.
- [31] O. A. LADYZHENSKAYA AND N. N. URAL'TSEVA, *Linear and Quasilinear Elliptic Equations*, Academic Press, New York, 1968.
- [32] J. C. MIELLOU, D. EL BAZ, AND P. SPITERI, *A new class of asynchronous iterative algorithms with order intervals*, Math. Comp., 67 (1998), pp. 237–255.
- [33] A. REUSKEN, *Convergence of the multigrid full approximation scheme for a class of elliptic mildly nonlinear boundary value problems*, Numer. Math., 52 (1988), pp. 251–277.
- [34] P. SPITERI, *Parallel asynchronous algorithms: a survey*, Adv. Eng. Softw., 149 (2020), Art. 102896, 34 pages.
- [35] P. SPITERI, J.-C. MIELLOU, AND D. EL BAZ, *Parallel asynchronous Schwarz and multisplitting methods for a nonlinear diffusion problem*, Numer. Algorithms, 33 (2003), pp. 461–474.
- [36] D. B. SZYLD, *Different models of parallel asynchronous iterations with overlapping blocks*, Comput. Appl. Math., 17 (1998), pp. 101–115.
- [37] J.-J. XU, *Convergence of partially asynchronous block quasi-Newton methods for nonlinear systems of equations*, J. Comput. Appl. Math., 103 (1999), pp. 307–321.