

**Asynchronous Domain  
Decomposition Methods  
for Nonlinear PDEs**

Fayçal Chaouqui, Edmond T. Chow  
and Daniel B. Szyld

Report 21-08-16  
August 2021. Revised July 2022

Department of Mathematics  
Temple University  
Philadelphia, PA 19122

This report is available in the World Wide Web at  
<http://www.math.temple.edu/~szyld>



1                   **ASYNCHRONOUS DOMAIN DECOMPOSITION METHODS**  
2                   **FOR NONLINEAR PDES\***

3                   FAYÇAL CHAOUQUI<sup>†</sup>, EDMOND CHOW<sup>‡</sup>, AND DANIEL B. SZYLD<sup>§</sup>

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

10                  **Key words.** Asynchronous iterations. Nonlinear problems. Domain decomposition. Partial  
11 Differential Equations. Two-level methods.

12                  **AMS subject classifications.** 65N22, 65F08, 65F10, 65F50

13                  **1. Introduction.** Asynchronous iterative methods are currently undergoing a  
14 resurgence in popularity due to the dramatic increase of parallelism in modern com-  
15 puters. The good performance of asynchronous methods is due to not needing to  
16 synchronize the computational tasks, hence minimizing idle time, i.e., time for which  
17 some processors are inactive but can be used. The effect of data exchange in the asyn-  
18 chronous method is then less pronounced; see, e.g., [2, 4, 6, 22]. In these methods,  
19 the iterations (or updates) are carried out in parallel by processors in an arbitrary  
20 order and without any synchronization. These asynchronous iterative methods are  
21 especially attractive when the network of processors is heterogeneous, or the commu-  
22 nication costs are large, or when the differences processors have different loads, e.g.,  
23 when each processor solves local problems with different physical properties or with  
24 different type of nonlinearities.

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

33                  Our point of departure is the restricted additive Schwarz (RAS) iterative method  
34 for linear systems [12], which we review in Section 2. We then introduce a restricted  
35 additive Schwarz method for nonlinear problems, and analyze its local convergence  
36 properties. We note that RAS has been used as preconditioners for the linear systems  
37 arising in Newton’s methods for nonlinear problems, e.g., in [10, 11, 17, 30]. In  
38 these methods, the linear systems at each step are solved using a Krylov subspace  
39 method and preconditioned with RAS. Since we have in mind parallel computers  
40 with high communication costs, the use of Krylov methods may not be suitable, since  
41 each iteration usually entails orthogonalizations with the concomitant inner products,

---

\*This version dated July 23, 2022. This research is supported in part by the U.S. Department of Energy under grant DE-SC0016578.

<sup>†</sup>Department of Mathematics, Temple University, Philadelphia, PA, USA, chaouqui@temple.edu

<sup>‡</sup>College of Computing, Georgia Institute of Technology, Atlanta, GA, USA, echow@cc.gatech.edu

<sup>§</sup>Department of Mathematics, Temple University, Philadelphia, PA, USA, szyld@temple.edu

42 which entail communication among all processors. This is why we concentrate on RAS  
 43 as a solver, not as a preconditioner. We note that earlier work has been done on using  
 44 RAS as a preconditioner; cf. [17, 13]. See also the recent work [29] for two-level  
 45 RAS preconditioners using additive/multiplicative coupling of coarse spaces, and [15]  
 46 for substructured coarse correction, and [24] for asynchronous multiplicative coarse  
 47 correction.

48 Our contribution includes the introduction of the asynchronous version of the  
 49 nonlinear restricted additive Schwarz method in Section 4, where we also state con-  
 50 vergence results, and discuss its implementation. We further introduce a coarse grid  
 51 correction in Section 5, and discuss convergence results of the two-level method.  
 52 Throughout the paper we illustrate the performance of the methods with numeri-  
 53 cal experiments.

54 **2. Domain Decomposition Methods.** Our goal is to study the parallel solu-  
 55 tion of discretized general nonlinear elliptic equations in  $\mathbb{R}^d$  ( $d \geq 2$ ) of the form

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

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

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

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

68 The problem (2.1) after discretization with, e.g., finite differences or finite ele-  
 69 ments, yields a large nonlinear system of equations of the form

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

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

74 **2.1. Restricted Additive Schwarz (RAS) for linear systems.** We intro-  
 75 duce in the next section the restricted Additive Schwarz (RAS) method for nonlinear  
 76 equations. We begin here by reviewing RAS for the case that  $F(u)$  is linear. Let us  
 77 thus consider

$$F(u) = Au - b, \quad (2.3)$$

78 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  $A$   
 79 is nonsingular, or in other words,  $A$  is of full rank, i.e.,  $\text{rank } A = n$ . The restricted

80 additive Schwarz introduced in [12] provides a means of constructing parallel iterative  
 81 solvers based on domain decomposition techniques. The general philosophy of RAS  
 82 is that the local problems (i.e., in each subdomain) are solved with overlap, but  
 83 the variables communicated to the other local problems are those corresponding to  
 84 subdomains without overlap.

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

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

87 These correspond to the variables of the subdomains without overlap. For the overlap-  
 88 ping subdomains, for any  $\delta > 0$ , we define  $S_i \subset S_{i,\delta} \subset S$ ,  $i = 1, \dots, P$ , thus defining  
 89 a covering of  $S$ . We set  $V_{i,\delta} := \mathbb{R}^{n_{i,\delta}}$  where  $n_{i,\delta} = |S_{i,\delta}|$ , the cardinality of the set  
 90 of variables corresponding to the  $i$ th subdomain with overlap  $\delta$ . We define also the  
 91 restriction mappings  $R_{i,\delta}: V \mapsto V_{i,\delta}$ ,  $i = 1, \dots, P$ . These restriction operators satisfy  
 92 the identities

$$R_{i,\delta} R_{i,\delta}^\top = I_{V_{i,\delta}}$$

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

93 Moreover, we have

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

94 which is known as the partition of unity condition.

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

$$u^{k+1} = u^k + \sum_{i=1}^P R_{i,0}^\top A_{i,\delta}^{-1} R_{i,\delta} (b - Au^k), \quad (2.4)$$

97 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 prob-  
 98 lem. Observe in (2.4) how the local problem is solved with the overlap, but only the  
 99 variables without the overlap contribute to the next iterate. The iteration operator  
 100 corresponding to (2.4) is thus given by

$$T_{A,\delta} = I - \sum_{i=1}^P R_{i,0}^\top A_{i,\delta}^{-1} R_{i,\delta} A. \quad (2.5)$$

101 We present a convergence and comparison result for (linear) RAS from [23]. We  
 102 need some notation first. Let  $\geq$  ( $>$ ) in  $\mathbb{R}^n$  and  $\mathbb{R}^{n \times n}$  denote the natural elementwise  
 103 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  $j =$   
 104  $1, \dots, n$ . Let  $w \in \mathbb{R}^n$ ,  $w > 0$ , and let  $\|\cdot\|_w$  denote the weighted max-norm operator  
 105 corresponding to the vector norm  $\|x\|_w = \max_{j=1, \dots, n} |x_j/w_j|$ . A nonsingular matrix  
 106  $A$  is called an  $M$ -matrix if its off-diagonal elements are nonpositive and the inverse is  
 107 nonnegative, i.e.,  $A^{-1} \geq 0$ .

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

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

110 Theorem 2.1 shows that the (linear) RAS iterative method case converges linearly,  
 111 and that the larger the overlap, the faster the asymptotic convergence. Recall that  
 112 the classical additive Schwarz method may fail to converge as an iterative method.  
 113 The main difference is that here the variables of the overlap are not included in the  
 114 new iterate (and thus there is no “double counting”); see [19, 23].

115 For completeness, we remark that when RAS is used as a preconditioner for  
 116 Krylov subspace methods, although this is not the focus of this paper, it follows from  
 117 (2.5) that the action of the RAS preconditioner  $M^{-1}$  on the global matrix  $A$  is given  
 118 by

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

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

122 **2.2. Nonlinear Restricted Additive Schwarz (NLRAS).** We return to the  
 123 case of nonlinear  $F$  and now describe the nonlinear iterative RAS (NLRAS) algorithm.  
 124 While this algorithm is new, we base our presentation on [18], where general Schwarz  
 125 methods for nonlinear problems are discussed. Define the restriction operators

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

126 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  
 127 nonlinear problem

$$F_{i,\delta}(u + R_{i,\delta}^\top v_{i,\delta}) = 0, \quad \text{for } i = 1, \dots, P. \quad (2.6)$$

128 We can then define the NLRAS fixed point iteration as

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

129 Note here again the philosophy of RAS: we solve the local (nonlinear) problem (2.6)  
 130 with the overlap, but only consider the non-overlapping variables as a contribution to  
 131 the next iterate.

132 We provide now a result showing that (2.7) converges locally and has an asymp-  
 133 totic convergence rate similar to that of the linear case. To this end, let  $T$  define the  
 134 nonlinear mapping

$$T(u) = u + \sum_{i=1}^P R_{i,0}^\top v_{i,\delta}, \quad (2.8)$$

135 where  $v_{i,\delta}$  is the solution of the local nonlinear problem (2.6). Let  $u^*$  be the solution  
 136 of (2.2). We shall prove that  $T$  is a contraction. Hence it admits a fixed point and  
 137 the iterative process (2.7) converges to  $u^*$ ; cf. [16].

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

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

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

142 *Proof.* By definition, we have

$$\begin{aligned} T(u) - T(u') &= \left( u + \sum_{i=1}^P R_{i,0}^\top v_{i,\delta} \right) - \left( u' + \sum_{i=1}^P R_{i,0}^\top v'_{i,\delta} \right) \\ &= u - u' + \sum_{i=1}^P R_{i,0}^\top (v_{i,\delta} - v'_{i,\delta}). \end{aligned} \quad (2.9)$$

143 Let us define the multidimensional differential quotient of  $F$  by

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

144 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  
 145  $v'_{i,\delta}$  satisfy

$$v_{i,\delta} - v'_{i,\delta} = -(R_{i,\delta} J_{i,\delta} R_{i,\delta}^\top)^{-1} R_{i,\delta} J_{i,\delta} (u - u'). \quad (2.10)$$

146 Indeed, using (2.6), we have that  $v_{i,\delta}$  and  $v'_{i,\delta}$  satisfy

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

147 from which we obtain,

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

148 Hence, substituting (2.10) in (2.9), we obtain

$$\begin{aligned} T(u) - T(u') &= u - u' - \sum_{i=1}^P R_{i,0}^\top (R_{i,\delta} J_{i,\delta} R_{i,\delta}^\top)^{-1} R_{i,\delta} J_{i,\delta} (u - u') \\ &= \left( I - \sum_{i=1}^P R_{i,0}^\top (R_{i,\delta} J_{i,\delta} R_{i,\delta}^\top)^{-1} R_{i,\delta} J_{i,\delta} \right) (u - u'). \end{aligned} \quad (2.11)$$

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

$$(R_{i,\delta} DF(u, u') R_{i,\delta}^\top)^{-1} R_{i,\delta} DF(u, u') \rightarrow (R_{i,\delta} F'(u^*) R_{i,\delta}^\top)^{-1} R_{i,\delta} F'(u^*) \text{ as } u, u' \rightarrow u^*. \quad (2.12)$$

150 Moreover, formula (2.10) shows that  $v_{i,\delta}, v'_{i,\delta} \rightarrow 0$ , when  $u, u' \rightarrow u^*$ . From this it  
 151 follows that

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

152 where we used the notation in (2.5). This shows that for  $\|T_{F'(u^*),\delta}\|_w < \zeta < 1$ , there  
 153 exists a neighborhood  $\mathcal{U}$  of  $u^*$  such that

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

154 It suffices then to take the norm  $\|\cdot\|_w$  on both sides of (2.11) to obtain the stated  
 155 result.  $\square$

156 Theorem 2.2 shows that the iterative procedure  $u^{k+1} = T(u^k)$  for  $u^0 \in \mathcal{U}$  con-  
 157 verges to  $u = T(u)$  for some  $u$ . This implies that we converge to the point  $u = u^*$ ,  
 158 the unique solution of (2.2). Moreover, the same result shows that the mapping  $T$  in  
 159 (2.8) is differentiable on  $u = u^*$ , with  $T'(u^*) = T_{F'(u^*),\delta}$ . This also demonstrates that  
 160 NLRAS with an overlap  $\delta' \geq \delta$  has in general a faster local convergence rate since  
 161  $\|T_{F'(u^*),\delta'}\|_w \leq \|T_{F'(u^*),\delta}\|_w$  (see Theorem 2.1).

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

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

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

170 is solved. However, we do not pursue this idea here, and we refer the interested reader  
 171 to, e.g., [10, 11, 17, 30].

172 **3. General Asynchronous iterations.** Asynchronous iterations refer to a class  
 173 of parallel iterative procedures in which each processor executes its computations at  
 174 the next iteration without waiting for the others to finish theirs. That is, there is  
 175 no synchronization. The concept of iteration thus loses its meaning, since different  
 176 processors update elements of the global approximation at different times (these are  
 177 called timestamps below). Thus some processors use information which may have  
 178 been updated in some processors more times than in others.

179 Mathematical models of asynchronous iterations were developed in order to study  
 180 their convergence; see, e.g., [2, 4, 6, 22, 34, 36]. Sometimes these models are referred  
 181 to as mathematical descriptions of computational models of asynchronous iterations.  
 182 We review one such model from [4], which is now classic. Let  $U = U_1 \times \dots \times U_P$ ,  
 183  $T : U \mapsto U$ , and  $T_i(u) = (Tu)_i$ . The goal is to solve (2.2) by way of the equivalent  
 184 fixed point problem  $u = T(u)$  in parallel, so that processor  $i$  runs iterations on the  
 185  $i$ th local problem, i.e.,  $u_i^{new} = T_i(u)$ . Let us call a *timestamp* the instant of time  
 186 at which at least one processor finishes its computation and updates its associated  
 187 local variables. Define by  $t_k$  the sequence of all these timestamps, and by  $s_j^i(k)$  the  
 188 sequence of integers that represent the timestamp index of the local variables coming  
 189 from processor  $j$ , and that are available to processor  $i$  at time instant when they start  
 190 solving the local problem which is completed at time  $t_k$ . Let us also denote by  $I^k$   
 191 a subset of  $\{1, \dots, P\}$  that defines the list of subdomains that are being updated at  
 192 timestamp  $t_k$ .



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

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

195 that is, either the  $i$ th portion of the solution is updated with the results of the compu-  
 196 tations in the  $i$ th processor at the timestamp  $t_k$ , or it is not updated. We also make  
 197 the following three assumptions.

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

201 The first assumption makes sure that no future updates are being used at the be-  
 202 ginning of the computation. The second ensures that eventually each processor is  
 203 receiving new information. The last one indicates that no processor stops being up-  
 204 dated. Convergence results for (3.1) were discussed and analyzed in several papers;  
 205 see, e.g., [2, 22, 34] and references therein.

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

208 **THEOREM 3.1.** ([22, Theorem 4.1] following [4]) *Let  $T$  be an  $n \times n$  matrix, i.e.,*  
 209 *representing a linear map. Assume that the conditions (i)–(iii) hold. If the spectral*  
 210 *radius  $\rho(|T|) < 1$ , where  $|T|$  has components  $|t_{ij}|$ , then the asynchronous iteration*  
 211 *(3.1) converges to the solution of the fixed point problem  $u = Tu$ .*

212 For example, when  $T$  is the RAS operator (2.5) and  $A$  is a nonsingular  $M$ -matrix,  
 213 we know from Theorem 2.1 that there exists a vector  $w > 0$  such that  $\|T_{A,\delta}\|_w < 1$ .  
 214 This implies that  $\rho(|T_{A,\delta}|) < 1$ , and we can use Theorem 3.1 to show that asyn-  
 215 chronous RAS converges.

216 **THEOREM 3.2.** ([22, Theorem 4.4] following [20]) *Assume that  $u^*$ , the solution*  
 217 *of (2.2), i.e., the fixed point of  $T$ , lies in the interior of  $U$  and that  $T$  is Fréchet*  
 218 *differentiable at  $u^*$ . Assume that the conditions (i)–(iii) hold and that  $\rho(|T'(u^*)|) < 1$ ,*  
 219 *then there exists a neighborhood  $\mathcal{U}$  of  $u^*$ , such that the asynchronous iteration (3.1)*  
 220 *converges to  $u^*$  for any  $u^0 \in \mathcal{U}$ .*

221 **THEOREM 3.3.** ([22, Theorem 3.3] following [20]) *Let  $T^k$  be a sequence of (non-*  
 222 *linear) operators,  $T^k : U \rightarrow U$ , such that there exists a common fixed point  $u^* \in U$*   
 223 *for which  $T^k(u^*) = u^*$  for all  $k$ . Moreover, assume that there exist  $0 \leq \zeta < 1$ , and*  
 224  *$w \in \mathbb{R}^n, w > 0$  such that*

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

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

$$u_i^{k+1} = \begin{cases} T_i^k(\dots, u_j^{s_j^i(k)}, \dots) & \text{for } i \in I^k \\ u_i^k & \text{for } i \notin I^k, \end{cases} \quad (3.2)$$

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

228 **4. Asynchronous nonlinear RAS.** Asynchronous parallel methods for non-  
 229 linear equations have been studied before; see, e.g. [3, 7, 14, 21, 25, 32, 35, 37]. The  
 230 experimental results obtained in these papers showed that the asynchronous methods

231 perform better in terms of execution times than their synchronous counterparts; see  
 232 e.g., [14, Tables C5–C7]. Here we present the asynchronous RAS method for nonlinear  
 233 problems, provide local convergence results, and compare its performance with that  
 234 of the synchronous counterpart. We emphasize that our approach is well suited to  
 235 heterogeneous networks, and in general to parallel computers where synchronization  
 236 and communication are at a premium.

237 The asynchronous nonlinear RAS method consists of the standard asynchronous  
 238 iteration (3.1), where  $T_i$  is given by (2.6) and (2.8); see also Algorithm 1 below. We  
 239 begin with a local convergence theorem.

240 **THEOREM 4.1.** *Let  $u^*$  be the solution of (2.2). Suppose that  $F'$  exists on a neigh-*  
 241 *borhood of  $u^*$  and that  $F'(u^*)$  is a nonsingular  $M$ -matrix. Assume that the conditions*  
 242 *(i)–(iii) hold. Then, there exists a neighborhood  $\mathcal{U}$  of  $u^*$  such that asynchronous non-*  
 243 *linear RAS (NLRAS) defined by (3.1) with  $T_i$  is given by (2.6) and (2.8), converges*  
 244 *to  $u^*$  for any initial vector  $u^0 \in \mathcal{U}$ .*

245 *Proof.* The proof is essentially the same as that in Theorem 2.2. However, here  
 246 we need the additional hypothesis that  $T$  is differentiable on a region containing  $u^*$ .  
 247 We recall that the nonlinear RAS iteration is given by

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

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

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

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

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

252 hence,

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

253 Thus,

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

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

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

255 where the latter is exactly the same as in (2.5) but applied to  $F'(u^*)$ , hence  $\|T'(u^*)\|_w =$   
 256  $\|T_{F'(u^*),\delta}\|_w < 1$  for  $w > 0$  such that  $F'(u^*)w > 0$ . Using Theorem 3.2, with the obser-  
 257 vation that  $\rho(|T_{F'(u^*),\delta}|) \leq \|T_{F'(u^*),\delta}\|_w$ , we conclude that there exists a neighborhood  
 258 of  $\mathcal{U}$  of  $u^*$  such that the asynchronous nonlinear RAS is convergent.  $\square$

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

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

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

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

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

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

- 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 (2.6) 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:**  $u^*$ .
-

278

279

280

281

282

283

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 check that it is below the convergence threshold. This was always satisfied in our experiments.

284

285

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

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

286

287

288

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 to the following nonlinear system of equations

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

289

290

291

292

293

294

295

296

297

298

299

300

301

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 [27]. The local problems in (2.6) 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's iterations is smaller than  $10^{-10}$ , and the local Jacobian matrices arising from the inner Newton's 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's iterations.

302

303

304

305

306

307

308

309

310

311

312

313

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's solve twice for one processor, and 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 there is a large load imbalance, that is, if one or more processors need to perform more work than others.

314

315

316

317

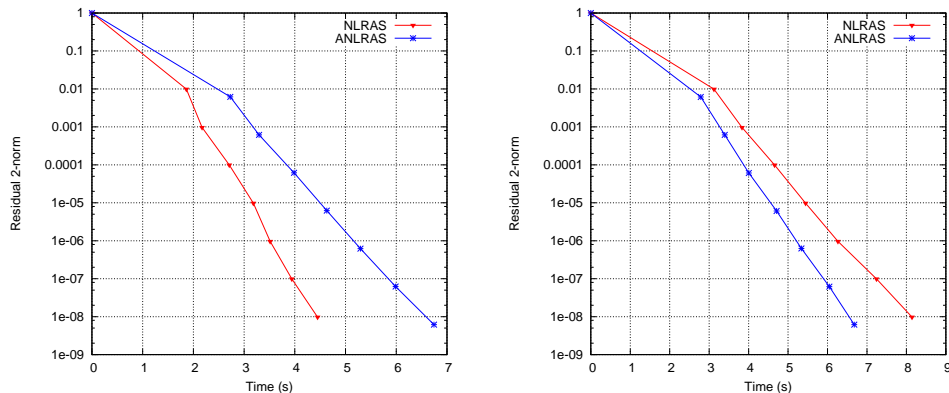
We report in Table 4.1 the number of local updates for different number of subdomains for both the synchronous and the asynchronous NLRAS. With one processor twice as slow, the asynchronous method is faster than the synchronous method even if, on average, the asynchronous method requires more local updates.

318

319

320

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. We can see from Figure 4.2 that the number of Newton iterations is slightly different



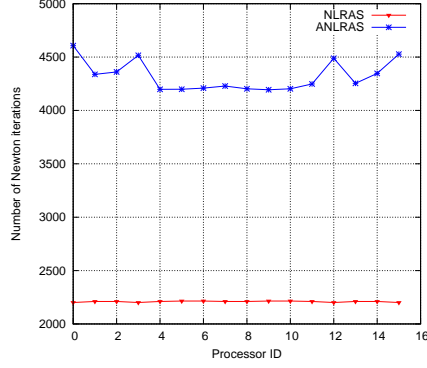
**Fig. 4.1:** Convergence curves of synchronous and asynchronous NLRAS for problem (4.4) with  $P = 16$ . Left: Homogeneous network. Right: One processor slower than the others by a factor of two.

#DOF	#subdomains	NLRAS	
		#iter (min,mean,max)	time
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

**Table 4.1:** Number of iterations/updates and computational time (in seconds) for problem (4.4) where one processor is twice as slow.

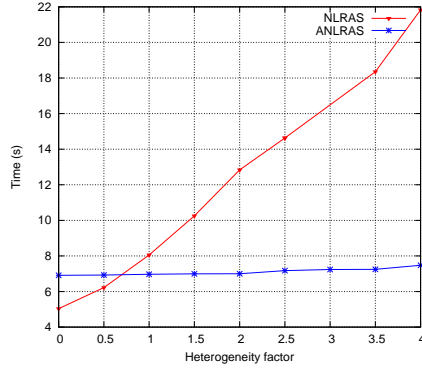
321 for each processor. We also observe that the asynchronous case requires more New-  
 322 ton iterations compared to the synchronous case since it usually requires more local  
 323 updates.

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



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

334 synchronous method increases roughly proportionally with the heterogeneity factor.




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

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

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

337 where  $v \in \mathbb{R}^2$ ,  $\Gamma \subset \partial\Omega$ , and  $\varphi: \mathbb{R} \mapsto \mathbb{R}$  is a function of the form  $\varphi(u) = \gamma \log(\beta + \alpha u)$ ,  
 338 with  $\alpha, \beta, \gamma > 0$ . For our experiments we consider  $\Omega = (0, 1)^2$ ,  $\Gamma = \{0\} \times [0, 1]$ , and

339  $v^T = [1, 1]$ . We discretize (4.6) using finite differences with  $n = 10000$  discretization  
 340 points. The resulting convergence plots are presented in Figure 4.4, and they illus-  
 341 trate the convergence of synchronous and asynchronous NLRAS for (4.6) for  $P = 16$   
 subdomains. We can clearly see that in this case the asynchronous method is more



**Fig. 4.4:** Convergence curves of synchronous and asynchronous NL-  
 RAS for problem (4.6) with  $\alpha = 1, \beta = 1, \gamma = 10^2$ . Left: Homogeneous  
 processor speed. Right: One processor is twice as slow.

342  
 343 advantageous, even in the case of an homogeneous network. This is due to the load  
 344 imbalance among subdomains solves. Some subdomains (those lying in the bound-  
 345 ary of  $\Omega$ ) require a nonlinear solver, while others require only a linear solver. The  
 346 advantage is enhanced when one of the processors is twice as slow as the others.

347 **5. Two-level nonlinear RAS (2L-NLRAS).** In order to improve the con-  
 348 vergence of NLRAS and ensure the scalability of the method, an additional coarse  
 349 correction can be incorporated. Since the problem is nonlinear, we need to use a non-  
 350 linear coarse grid correction. We opt here for the FAS (Full Approximation Scheme)  
 351 method to perform the coarse grid correction; see, e.g., [9]. Let  $V_0$  denote the coarse  
 352 space and  $R_0 : V \mapsto V_0$  the coarse restriction operator. We can define then by means  
 353 of  $R_0$  the coarse nonlinear problem  $F_0 : V_0 \mapsto V_0$  which could be defined by using a  
 354 coarse discretization of (2.1), or using a Galerkin approach, namely

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

355 The coarse correction  $v_0 \in V_0$  using the Full Approximation Scheme is defined as the  
 356 solution of the nonlinear problem

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

357 which can be incorporated into (2.7) either in an additive or multiplicative manner.  
 358 Following [26], here we opt to use the additive variant. In order to prevent overcor-  
 359 rection on the overlap, we use a weighted two-level method of the form

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

see [24] for an alternative approach. The Full Approximation Scheme (FAS) [8] is a particular case of the nonlinear multilevel method (NMLM) of Hackbusch [28]. Although NMLM is a globally convergent iteration [28], 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.,

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

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

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

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

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

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

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

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 [5, Theorem 7.3] to complete the proof, as we show below. Let

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

and

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



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

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

388 and

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

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

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

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

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

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

393 We remark that  $B$  and  $B_0$  are restrictions and prolongations of nonnegative ma-  
 394 trices, where the restrictions and prolongation operators are nonnegative, they are  
 395 nonnegative. Furthermore,  $B \geq 0$  is nonsingular, in fact, it is a RAS preconditioner.  
 396 Therefore, it cannot have a zero row, and thus  $Be > 0$ , for any positive vector  $e$ . Let  
 397  $e = F'(u^*)w > 0$ . Then, we have that  $Be > 0$  and  $B_0e \geq 0$ , so that


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

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

399 We note that the hypotheses mentioned in Theorem (5.1) are not always satisfied.  
 400 In particular, for our coarse grid, the nonnegativity of the operator  $\tilde{T}'$  does not nec-  
 401 essarily hold. For the particular case where the coarse grid corresponds to a boolean  
 402 matrix, it is shown in [1] that the choice of  $E_0 = R_0^\top R_0$  is possible.

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

406 We show in Figure 5.1 a weak scaling experiment for the problem (4.4). We com-  
 407 pare the one- and two-level NLRAS for a  $2 \times 2$  and  $3 \times 3$  decomposition. The total  
 408 number of degrees of freedom (DOFs) in each subdomain is kept the same for both  
 409 cases, namely 625 variables per subdomain. As we can observe, the convergence of the  
 410 one-level NLRAS deteriorates when the number of subdomains increases from  $P = 4$   
 411 to  $P = 9$ . In contrast, the convergence rate of the two-level NLRAS remains approx-  
 412 imately the same. We summarize weak scalability results with increasing numbers of  
 413 subdomains for iteration (5.2) in Table 5.1, where we report number of iterations to  
 414 reach the tolerance of  $10^{-8}$ . These results show that the method without coarse cor-  
 415 rection fails to converge in 100 iterations, while for the two-level NLRAS, the number  
 416 of iterations has a slow growth with the number of subdomains.



**Fig. 5.1:** Weak scaling experiment for one- and two-level NLRAS for problem (4.4) using a  $2 \times 2$  checker-board partition (left), and  $3 \times 3$  checker-board partition (right)

$P$	$ 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

**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  $\epsilon = 10^{-8}$

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

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

419 We summarize in Table 5.2 the convergence results of iteration (5.2) for this problem.  
 420 We point out that for a discretization of (5.6) using finite differences,  $F'(u)$  is a  
 421 nonsingular  $M$ -matrix for  $u \in \Omega$ , and therefore the hypothesis of our local convergence  
 422 theorems are satisfied in this case. The results in Table 5.2 are similar to those in  
 423 Table 5.1. The addition of the coarse correction makes the algorithm scalable in the  
 424 sense that the number of iterations required to achieve the prescribed tolerance has  
 425 a very slow growth as the number of subdomains increases. Comparing Tables 5.1  
 426 and 5.2, we note that for two problems with different type of nonlinearity, the two-level  
 427 method behaves similarly in terms of number of iterations to convergence.

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

$P$	$ 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

**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  $\epsilon = 10^{-8}$

432 as follows. By recalling the definition of  $F_{i,\delta}$ , we have that

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

433 Similarly, we obtain for the derivative of  $F_0$

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

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

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

439 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  
440 nonlinear residual  $R_{i,\delta} F(u^k)$  is zero outside the overlap. This shows that a good coarse  
441 grid should use this information for the construction of the coarse space. Thus, the  
442 coarse grid should be chosen on the overlap.

443 **6. Asynchronous Two-Level NLRAS.** In this section, we study the asyn-  
444 chronous version of the iteration (5.2). We present implementation details, local  
445 convergence theory, and numerical experiments. The information for the coarse grid  
446 correction comes from all subdomains, but it cannot be used until the information  
447 from all subdomains has arrived. Only then, the correction is added, and the process  
448 repeats. We implement this following ideas described in [26] for linear problems. We  
449 use a weighted additive NLRAS method (see [24] for multiplicative variants), and in  
450 order to prevent the coarse correction from over-correcting, we add the coarse grid  
451 only when all the subdomains have sent new data. This can be realized by defining  
452 a local boolean variable for each thread/subdomain. The coarse grid is then added  
453 when all these local boolean variables are set to true and this is done without a syn-  
454 chronization point. We describe the asynchronous two-level NLRAS in Algorithm 2.

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

---

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

---

**Require:** Initial vector  $u^0$ , tolerance  $\epsilon$ .

```

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 (2.6) 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:  $u^*$ .

```

---

457 **THEOREM 6.1.** *Let  $u^*$  be the solution of (2.2). Assume that  $F'$  exists on a*  
458 *neighborhood of  $u^*$  and that  $F'(u^*)$  and  $F'(R_0^\top R_0 u^*)$  are nonsingular  $M$ -matrices.*  
459 *Let  $w > 0$  be such that  $T'(u^*)w > 0$ . Assume that the conditions (i)–(iii) hold. Then*  
460 *Algorithm 2 converges to  $u^*$ , the solution of (2.2).*

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


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

463 if the coarse grid in Algorithm 2 is added, or

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

464 otherwise. Of course the operators (6.1) are those of the two-level NLRAS method (5.3),  
465 while those in (6.2) are those of the one-level method (2.8). It follows that both opera-


466 tors have  $u^*$  as a common fixed point, and in both cases,  $\|T^k\|_w < 1$ , by Theorems 2.2  
 467 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. .

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

483 We conclude with a table comparing execution times and the number of processor  
 484 updates for the one-level and two-level methods in both the synchronous and asyn-  
 485 chronous versions for problem (4.4) in 2D and 3D using  $P + 1$  processors. We note  
 486 that each subdomain is assigned to a processor and the coarse grid is computed on a  
 487 separate one. In the 3D case, the domain  $\Omega = (0, 1)^3$  is decomposed into  $P$  equally  
 488 sized bricks, and their number is the same in each direction. The size of each local  
 489 problem is  $1000 \times 1000$ . We assume that one processor is four times slower than the  
 490 others. The convergence results are reported in Table 6.1. We see from Table 6.1  
 491 that the asynchronous method outperforms the synchronous NLRAS. The addition  
 492 of the coarse grid correction makes both methods even faster, with an advantage for



**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 other ones. Left: Problem (4.4). Right: Problem (5.6).

493 the asynchronous one. The results in Table 6.1 also shows that the timing required  
 494 for the coarse step increases for increasing number of subdomains.

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

P	V   V <sub>0</sub>		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	

**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  $\epsilon = 10^{-8}$  for different numbers of subdomains with constant local problem size.

498

499 **7. Conclusion.** In summary, the new parallel nonlinear RAS algorithm con-  
 500 verges to the unique solution of the problem in the original domain of the PDE. An  
 501 asynchronous version can be about than twice as fast or faster, depending on the  
 502 inhomogeneities of the computation, and the communication costs. The two-level

503 version can be an order of magnitude faster, and its asynchronous version faster still,  
504 exhibiting slow growth in execution times when the number of processors increases in  
505 weak scaling.

- [1] Andreas Frommer and Daniel B. Szyld. Weighted Max Norms, Splittings, and Overlapping Additive Schwarz Iterations. *Numerische Mathematik*, 83:259–278, 1999.
- [2] Jacques Mohcine Bahi, Sylvain Contassot-Vivier, and Raphaël Couturier. *Parallel Iterative Algorithms: from Sequential to Grid Computing*. CRC Press, 2007.
- [3] Zhong-Zhi Bai, Violeta Migallón, José Penadés, and Daniel B. Szyld. Block and asynchronous two-stage methods for mildly nonlinear systems. *Numerische Mathematik*, 82:1–20, 1999.
- [4] Gérard M. Baudet. Asynchronous iterative methods for multiprocessors. *Journal of the Association for Computing Machinery*, 25:226–244, 1978.
- [5] Michele Benzi, Andreas Frommer, Reinhard Nabben, and Daniel B. Szyld. Algebraic theory of multiplicative Schwarz methods. *Numerische Mathematik*, 89(4):605–639, 2001.
- [6] Dimitri P. Bertsekas and John N. Tsitsiklis. *Parallel and distributed computation: numerical methods*, volume 23. Prentice hall Englewood Cliffs, NJ, 1989.
- [7] Amit Bhaya, Eugenius Kaszurewicz, and Francisco Mota. Asynchronous block-iterative methods for almost linear equations. *Linear Algebra and its Applications*, 154:487–508, 1991.
- [8] Achi Brandt. Multi-level adaptive solutions to boundary-value problems. *Mathematics of Computation*, 31(138):333–390, 1977.
- [9] Achi Brandt. Guide to multigrid development. In Wolfgang Hackbusch and Ulrich Trottenberg, editors, *Multigrid methods*, pages 220–312. Springer, 1982.
- [10] Xiao-Chuan Cai and David E. Keyes. Nonlinearly preconditioned inexact Newton algorithms. *SIAM Journal on Scientific Computing*, 24(1):183–200, 2002.
- [11] Xiao-Chuan Cai and Xuefeng Li. Inexact Newton methods with restricted additive Schwarz based nonlinear elimination for problems with high local nonlinearity. *SIAM Journal on Scientific Computing*, 33(2):746–762, 2011.
- [12] Xiao-Chuan Cai and Marcus Sarkis. A restricted additive Schwarz preconditioner for general sparse linear systems. *SIAM Journal on Scientific Computing*, 21(2):792–797, 1999.
- [13] Faycal Chaouqui, Martin J. Gander, Pratik M. Kumbhar, and Tommaso Vanzan. Linear and nonlinear substructured Restricted Additive Schwarz iterations and preconditioning. *Numerical algorithms*, inprint, 2022.
- [14] Ming Chau, Didier El Baz, Ronan Guivarch, and Pierre Spiteri. MPI implementation of parallel subdomain methods for linear and nonlinear convection–diffusion problems. *Journal of Parallel and Distributed Computing*, 67(5):581–591, 2007.
- [15] Gabriele Ciaramella and Tommaso Vanzan. Substructured two-grid and multi-grid domain decomposition methods. *Numerical Algorithms*, pages 1–36, 2022.
- [16] Ljubomir B. Ciric. Generalized contractions and fixed-point theorems. *Publications de l’Institut Mathématique (Beograd)(NS)*, 12(26):19–26, 1971.
- [17] Victorita Dolean, Martin J. Gander, Walid Kheriji, Felix Kwok, and Roland Masson. Nonlinear preconditioning: How to use a nonlinear Schwarz method to precondition Newton’s method. *SIAM Journal on Scientific Computing*, 38(6):A3357–A3380, 2016.
- [18] Maksymilian Dryja and Wolfgang Hackbusch. On the nonlinear domain decomposition method. *BIT Numerical Mathematics*, 37(2):296–311, 1997.
- [19] Evridiki Efstathiou and Martin J. Gander. Why restricted additive Schwarz converges faster than additive Schwarz. *BIT Numerical Mathematics*, 43(5):945–959, 2003.
- [20] Mouhamed Nabih El Tarazi. Some convergence results for asynchronous algorithms. *Numerische Mathematik*, 39(3):325–340, 1982.
- [21] Andreas Frommer and Hartmut Schwandt. Asynchronous parallel methods for enclosing solutions of nonlinear equations. *Journal of Computational and Applied Mathematics*, 60(1-2):47–62, 1995.
- [22] Andreas Frommer and Daniel B Szyld. On asynchronous iterations. *Journal of Computational and Applied Mathematics*, 123(1-2):201–216, 2000.
- [23] Andreas Frommer and Daniel B. Szyld. An algebraic convergence theory for restricted additive Schwarz methods using weighted max norms. *SIAM Journal on Numerical Analysis*, 39:463–479, 2001.
- [24] Guillaume Gbikpi Benissan and Frédéric Magoulès. Asynchronous Multiplicative Coarse-Space Correction. *SIAM Journal on Scientific Computing*, 44(3):C237–C259, 2022.
- [25] Luc Giraud and Pierre Spiteri. Résolution parallèle de problèmes aux limites non linéaires. *ESAIM: Mathematical Modelling and Numerical Analysis*, 25(5):579–606, 1991.
- [26] Christian Glusa, Erik G. Boman, Edmond Chow, Sivasankaran Rajamanickam, and Daniel B. Szyld. Scalable Asynchronous Domain Decomposition Solvers. *SIAM Journal on Scientific Computing*, pages C384–C409, 2020.
- [27] Gaël Guennebaud, Benoit Jacob, et al. Eigen. URL: <http://eigen.tuxfamily.org>, 2010.



- 567 [28] Wolfgang Hackbusch. *Multigrid Methods and Applications*, volume 4 of *Springer Series in*  
568 *Computational Mathematics*. Springer, Berlin Heidelberg, 1985.
- 569 [29] Alexander Heinlein and Martin Lanser. Additive and Hybrid Nonlinear Two-Level Schwarz  
570 Methods and Energy Minimizing Coarse Spaces for Unstructured Grids. *SIAM Journal on*  
571 *Scientific Computing*, 42(4):A2461–A2488, 2020.
- 572 [30] Feng-Nan Hwang and Xiao-Chuan Cai. Improving Robustness and Parallel Scalability of  
573 Newton Method Through Nonlinear Preconditioning. In Ralf Kornhuber, Ronald Hoppe,  
574 Jacques Périaux, Olivier Pironneau, Olof Widlund, and Jinchao Xu, editors, *Domain De-*  
575 *composition Methods in Science and Engineering*, volume 40 of *Lecture Notes in Compu-*  
576 *tational Science and Engineering*, pages 201–208, Berlin, Heidelberg, 2005. Springer.
- 577 [31] Olga A. Ladyzhenskaya and Nina N. Ural'tseva. *Linear and Quasilinear Elliptic Equations*.  
578 *Academic Press, New York*, 1968.
- 579 [32] Jean-Claude Miellou, Didier El Baz, and Pierre Spiteri. A new class of asynchronous iterative  
580 algorithms with order intervals. *Mathematics of Computation*, 67(221):237–255, 1998.
- 581 [33] Arnold Reusken. Convergence of the multigrid full approximation scheme for a class of elliptic  
582 mildly nonlinear boundary value problems. *Numerische Mathematik*, 52(3):251–277, 1987.
- 583 [34] Pierre Spiteri. Parallel asynchronous algorithms: A survey. *Advances in Engineering Software*,  
584 149:102896, 2020.
- 585 [35] Pierre Spiteri, Jean-Claude Miellou, and Didier El Baz. Parallel asynchronous Schwarz and  
586 multisplitting methods for a nonlinear diffusion problem. *Numerical Algorithms*, 33(1-  
587 4):461–474, 2003.
- 588 [36] Daniel B. Szyld. Different models of parallel asynchronous iterations with overlapping blocks.  
589 *Computational and Applied Mathematics*, 17:101–115, 1998.
- 590 [37] Jian-Jun Xu. Convergence of partially asynchronous block quasi-Newton methods for nonlinear  
591 systems of equations. *Journal of Computational and Applied Mathematics*, 103(2):307–321,  
592 1999.