

**Asynchronous Domain  
Decomposition Methods  
for Nonlinear PDEs**

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

Report 21-08-16  
August 2021

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 T. 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., [1, 3, 5, 20]. 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 [11], 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 [9, 10, 15, 26]. In these  
38 methods, the linear systems at each step are solved using a Krylov subspace method  
39 and preconditioned with RAS. Since we have in mind parallel computers with high  
40 communication costs, the use of Krylov methods may not be suitable, since each iter-  
41 ation usually entails orthogonalizations with the concomitant inner products, which

---

\*This version dated August 15, 2021. 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 entail communication among all processors. This is why we concentrate on RAS as a  
 43 solver, not as a preconditioner. We believe that this is the first time that a nonlinear  
 44 version of RAS is proposed.

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

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

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

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

- 59 (i)  $a_i \in C^1(\Omega \times \mathbb{R}^{d+1})$
- 60 (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  $j = 1, \dots, d$
- 61 (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  $(\xi_0, \dots, \xi_d) \in \mathbb{R}^{d+1}$ .

62 Then, the problem (2.1) is well posed in  $H^1(\Omega)$ , where  $H^1(\Omega)$  is the space of func-  
 63 tions  $v \in L^2(\Omega)$  such that  $\nabla v \in L^2(\Omega)$ , and the space is equipped with the norm  
 64  $\|v\|_{H^1(\Omega)}^2 := \|v\|_{L^2(\Omega)}^2 + \|\nabla v\|_{L^2(\Omega)}^2$ ; see [27].

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

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

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

71 **2.1. Restricted Additive Schwarz (RAS) for linear systems.** We intro-  
 72 duce in the next section the restricted Additive Schwarz (RAS) method for nonlinear  
 73 equations. We begin here by reviewing RAS for the case that  $F(u)$  is linear. Let us  
 74 thus consider

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

75 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  $A$   
 76 is nonsingular, or in other words,  $A$  is of full rank, i.e.,  $\text{rank } A = n$ . The restricted  
 77 additive Schwarz introduced in [11] provides a means of constructing parallel iterative  
 78 solvers based on domain decomposition techniques. The general philosophy of RAS  
 79 is that the local problems (i.e., in each subdomain) are solved with overlap, but

80 the variables communicated to the other local problems are those corresponding to  
81 subdomains without overlap.

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

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

84 These correspond to the variables of the subdomains without overlap. For the overlap-  
85 ping subdomains, for any  $\delta > 0$ , we define  $S_i \subset S_{i,\delta} \subset S$ ,  $i = 1, \dots, P$ , thus defining  
86 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  
87 of variables corresponding to the  $i$ th subdomain with overlap  $\delta$ . We define also the  
88 restriction mappings  $R_{i,\delta}: V \mapsto V_{i,\delta}$ ,  $i = 1, \dots, P$ . These restriction operators satisfy  
89 the identities

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

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

90 Moreover, we have

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

91 which is known as the partition of unity condition.

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

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

94 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-  
95 lem. Observe in (2.4) how the local problem is solved with the overlap, but only the  
96 variables without the overlap contribute to the next iterate. The iteration operator  
97 corresponding to (2.4) is thus given by

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

98 We present a convergence and comparison result for (linear) RAS from [21]. We  
99 need some notation first. Let  $\geq$  ( $>$ ) in  $\mathbb{R}^n$  and  $\mathbb{R}^{n \times n}$  denote the natural elementwise  
100 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  $j =$   
101  $1, \dots, n$ . Let  $w \in \mathbb{R}^n$ ,  $w > 0$ , and let  $\|\cdot\|_w$  denote the weighted max-norm operator  
102 corresponding to the vector norm  $\|x\|_w = \max_{j=1, \dots, n} |x_j/w_j|$ . A nonsingular matrix  
103  $A$  is called an  $M$ -matrix if its off-diagonal elements are nonpositive and the inverse is  
104 nonnegative, i.e.,  $A^{-1} \geq 0$ .

105 **THEOREM 2.1.** *Let  $A$  be a nonsingular  $M$ -matrix. Let  $w > 0$  such that  $Aw > 0$ .  
106 Then, if  $\delta \geq \delta'$ ,*

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

107 Theorem 2.1 shows that the (linear) RAS iterative method case converges linearly,  
 108 and that the larger the overlap, the faster the asymptotic convergence. Recall that  
 109 the classical additive Schwarz method may fail to converge as an iterative method.  
 110 The main difference is that here the variables of the overlap are not included in the  
 111 new iterate (and thus there is no “double counting”); see [17, 21].

112 For completeness, we remark that when RAS is used as a preconditioner for  
 113 Krylov subspace methods, although this is not the focus of this paper, it follows from  
 114 (2.5) that the action of the RAS preconditioner  $M^{-1}$  on the global matrix  $A$  is given  
 115 by

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

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

119 **2.2. Nonlinear Restricted Additive Schwarz (NLRAS).** We return to the  
 120 case of nonlinear  $F$  and now describe the nonlinear iterative RAS (NLRAS) algo-  
 121 rithm. While this algorithm is new, we base our presentation on [16], where general  
 122 Schwarz methods for nonlinear problems are discussed. Define the restrictiond non-  
 123 linear operators

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

124 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  
 125 nonlinear problem

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

126 We can then define the NLRAS fixed point iteration as

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

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

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

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

133 where  $v_{i,\delta}$  is the solution of the local nonlinear problem (2.6). Let  $u^*$  be the solution  
 134 of (2.2). We shall prove that  $T$  is a contraction. Hence it admits a fixed point and  
 135 the iterative process (2.7) converges to  $u^*$ ; cf. [14].

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

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

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

140 *Proof.* By definition, we have

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

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

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

142 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  
 143  $v'_{i,\delta}$  satisfy

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

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

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

145 from which we obtain,

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

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

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

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

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

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

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

150 where we used the notation in (2.5). Therefore for  $\delta' > 0$  fixed and  $\delta > \delta'$ , using  
 151 Theorem 2.1, we get

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

152 in a neighborhood  $\mathcal{U}$  of  $u^*$ . It suffices then to take the norm  $\|\cdot\|_w$  on both sides  
 153 of (2.11) to obtain the stated result.  $\square$

154 Theorem 2.2 shows that the iterative procedure  $u^{k+1} = T(u^k)$  for  $u^0 \in \mathcal{U}$  con-  
 155 verges to  $u = T(u)$  for some  $u$ . This implies that we converge to the point  $u = u^*$ ,  
 156 the unique solution of (2.2). Moreover, the same result shows that the mapping  $T$  in  
 157 (2.8) is differentiable on  $u = u^*$ , with  $T'(u^*) = T_{F'(u^*),\delta}$ .

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

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

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

166 is solved. However, we do not pursue this idea here, and we refer the interested reader  
 167 to, e.g., [9, 10, 15, 26].

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

175 Mathematical models of asynchronous iterations were developed in order to study  
 176 their convergence; see, e.g., [1, 3, 5, 20, 30, ?]. Sometimes these models are referred  
 177 to as mathematical descriptions of computational models of asynchronous iterations.  
 178 We review one such model from [3], which is now classic. Let  $U = U_1 \times \dots \times U_p$ ,  
 179  $T : U \mapsto U$ , and  $T_i(u) = (Tu)_i$ . The goal is to solve (2.2) by way of the equivalent  
 180 fixed point problem  $u = T(u)$  in parallel, so that processor  $i$  runs iterations on the  
 181  $i$ th local problem, i.e.,  $u_i^{new} = T_i(u)$ . Let us call a *timestamp* the instant of time  
 182 at which at least one processor finishes its computation and updates its associated  
 183 local variables. Define by  $t_k$  the sequence of all these timestamps, and by  $s_j^i(k)$  the  
 184 sequence of integers that represent the timestamp index of the local variables coming  
 185 from processor  $j$ , and that are available to processor  $i$  at time instant when they  
 186 start solving the local problem which is completed at time  $t_k$ . Let us also denote by  
 187  $I^k$  a subset of  $\{1, \dots, P\}$  that defines the list of variables that are being updated at  
 188 timestamp  $t_k$ .

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



190  $i = 1, \dots, P$ , satisfying

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

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

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

197 The first assumption makes sure that no future updates are being used at the be-  
 198 ginning of the computation. The second ensures that eventually each processor is  
 199 receiving new information. The last one indicates that no processor stops being up-  
 200 dated. Convergence results for (3.1) were discussed and analyzed in several papers;  
 201 see, e.g., [1, 20, 30] and references therein.

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

204 **THEOREM 3.1.** ([20, Theorem 4.1] following [3]) *Let  $T$  be an  $n \times n$  matrix, i.e.,*  
 205 *representing a linear map. Assume that the conditions (i)–(iii) hold. If the spectral*  
 206 *radius  $\rho(|T|) < 1$ , where  $|T|$  has components  $|t_{ij}|$ , then the asynchronous iteration*  
 207 *(3.1) converges to the solution of the fixed point problem  $u = Tu$ .*

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

212 **THEOREM 3.2.** ([20, Theorem 4.4] following [18]) *Assume that  $u^*$ , the solution*  
 213 *of (2.2), i.e., the fixed point of  $T$ , lies in the interior of  $U$  and that  $T$  is Fréchet*  
 214 *differentiable at  $u^*$ . Assume that the conditions (i)–(iii) hold and that  $\rho(|T'(u^*)|) < 1$ ,*  
 215 *then there exists a neighborhood  $\mathcal{U}$  of  $u^*$ , such that the asynchronous iteration (3.1)*  
 216 *converges to  $u^*$  for any  $u^0 \in \mathcal{U}$ .*

217 **THEOREM 3.3.** ([20, Theorem 3.3] following [18]) *Let  $T^k$  be a sequence of (non-*  
 218 *linear) operators,  $T^k : U \rightarrow U$ , such that there exists a common fixed point  $u^* \in U$*   
 219 *for which  $T^k(u^*) = u^*$  for all  $k$ . Moreover, assume that there exist  $0 \leq \zeta < 1$ , and*  
 220  *$w \in \mathbb{R}^n, w > 0$  such that*

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

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

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

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

224 **4. Asynchronous nonlinear RAS.** Asynchronous parallel methods for non-  
 225 linear equations have been studied before; see, e.g. [2, 6, 12, 19, 22, 28, 31, 32]. The  
 226 experimental results obtained in these papers showed that the asynchronous methods  
 227 perform better in terms of execution times than their synchronous counterparts; see

228 e.g., [12, Tables C5–C7]. Here we present the asynchronous RAS method for nonlinear  
 229 problems, provide local convergence results, and compare its performance with that  
 230 of the synchronous counterpart. We emphasize that our approach is well suited to  
 231 heterogeneous networks, and in general to parallel computers where synchronization  
 232 and communication are at a premium.

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

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

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

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

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

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

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

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

248 hence,

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

249 Thus,

$$T'(u) = I - \sum_{i=1}^P R_{i,\delta}^\top (R_{i,\delta} F'(u + R_{i,\delta}^\top v_{i,\delta}) R_{i,\delta})^{-1} F'_{i,\delta}(u + R_{i,\delta}^\top v_{i,\delta}). \quad (4.1)$$

250 Evaluating (4.1) at  $u = u^*$ , we obtain

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

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

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

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

$$\|F(u^k)\|_2 / \|F(u^0)\|_2 < \epsilon. \quad (4.2)$$

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

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

265 Hence, in order to check that the global stopping criterion was achieved, it suffices  
 266 to check that the local solution on each subdomain satisfies the local convergence  
 267 criteria  $\|F_{i,0}(u)\|_2 / \|F(u^0)\|_2 < \epsilon / \sqrt{P}$ . This condition ensures that the contribution  
 268 of the local residual is such that (4.2) holds. In our implementation, a global boolean  
 269 variable is declared and is then set to true when all the subdomains satisfy the local  
 270 convergence criteria. 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:**  $u^*$ .
- 

271 We mention that once a processor/subdomain achieves local convergence, the  
 272 local convergence can be lost before the global convergence flag is triggered. This  
 273 is why we compute the global residual norm again after the iterations stopped and  
 274 check that it is below the convergence threshold. This was always satisfied in our  
 275 experiments.  
 276


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

$$-\nabla^2 u + g(u) = f \quad \text{on } \Omega, \quad u = 0 \quad \text{on } \partial\Omega, \quad (4.3)$$

279 where  $\Omega = (0, 1)^2$  and  $g(u) = ue^u$ . We choose  $f$  such that  $\sin(\pi x)\sin(\pi y)$  is the exact  
 280 solution of (4.3). The discretization of (4.3) using standard 5-points stencil with mesh  
 281 size  $h$  on the nodes  $(ih, jh)$  yields to the following nonlinear system of equations

$$(F(u))_{i,j} = \frac{1}{h^2}(4u_{i,j} - u_{i-1,j} - u_{i+1,j} - u_{i,j+1} - u_{i,j-1}) + u_{i,j}e^{u_{i,j}} - f_{i,j} = 0, \quad (4.4)$$

282 with  $u_{0,j} = u_{N+1,j} = u_{j,0} = u_{j,N+1} = 0$ . We present and analyze computational  
 283 results for parallel asynchronous and synchronous NLRAS for the discrete nonlin-  
 284 ear system of equations (4.4). The computational experiments were carried out on  
 285 a shared memory machine with 88 CPU cores/176 threads and 1536GB of RAM.  
 286 The parallelization of the solvers was implemented in C++ using the OpenMP mul-  
 287 tithreading library. The linear algebra data structures and solvers were provided by  
 288 EIGEN [24]. The local problems in (2.6) were solved using an inner Newton's method  
 289 where the local convergence criterion is when the norm of the difference between two  
 290 updates in the local Newton's iterations is smaller than  $10^{-10}$ , and the local Jacobian  
 291 matrices arising from the inner Newton's iterations were computed analytically. We  
 292 use a checker-board partitioning for the domain  $\Omega = (0, 1)^2$  with a total of 10000  
 293 discretization points. We choose the zero vector as an initial vector for both the RAS  
 294 iteration and the local Newton's iterations.



**Fig. 4.1:** Convergence curves of synchronous and asynchronous NLRAS for problem (4.3) with  $P = 16$ . Left: Homogeneous network. Right: One processor slower than the others by a factor of two.


295 In Figure 4.1 (left), we compare both the synchronous and asynchronous NLRAS  
 296 for a partition into 16 subdomains, and without any load imbalance. We plot the  
 297 timing versus the relative residual norm for both methods. We observe in this case  
 298 that the synchronous method is faster in terms of execution time. We repeat the same

299 experiment, but with one processor twice as slow as the others by performing the inner  
 300 Newton's solve twice for one processor, and report the result in Figure 4.1 (right).  
 301 It is clear that in this case, the asynchronous method outperforms the synchronous  
 302 one in terms of execution time. This indicates that even with a small number of  
 303 processors, if the network is heterogeneous, asynchronous methods may be preferred.  
 304 The same effect would be observed if there is a large load imbalance, that is, if one  
 305 or more processors need to perform more work than others.

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

$$\begin{aligned}
 -\nabla^2 u + v \cdot \nabla u &= f & \text{on } \Omega \\
 \frac{\partial u}{\partial \nu} + \varphi(u) &= 0 & \text{on } \Gamma \quad \text{and } u = 0 \text{ on } \partial\Omega \setminus \Gamma,
 \end{aligned}
 \tag{4.5}$$

308 where  $v \in \mathbb{R}^2$ ,  $\Gamma \subset \partial\Omega$ , and  $\varphi: \mathbb{R} \mapsto \mathbb{R}$  is a function of the form  $\varphi(u) = \gamma \log(\beta + \alpha u)$ ,  
 309 with  $\alpha, \beta, \gamma > 0$ . For our experiments we consider  $\Omega = (0, 1)^2$ ,  $\Gamma = \{0\} \times [0, 1]$ , and  
 310  $v^T = [1, 1]$ . We discretize (4.5) using finite differences with  $n = 10000$  discretization  
 311 points. The resulting convergence plots are presented in Figure 4.2, and they illus-  
 312 trate the convergence of synchronous and asynchronous NLRAS for (4.5) for  $P = 16$   
 subdomains. We can clearly see that in this case the asynchronous method is more



**Fig. 4.2:** Convergence curves of synchronous and asynchronous NL-  
 RAS for problem (4.5) with  $\alpha = 1, \beta = 1, \gamma = 10^2$ . Left: Homogeneous  
 processor speed. Right: One processor slower.

313 advantageous, even in the case of a homogeneous network. This is due to the load im-  
 314 balance among subdomains solves. Some subdomains (those lying in the boundary of  
 315  $\Omega$ ) require a nonlinear solver, while others require only a linear solver. The advantage  
 316 is enhanced when one of the processors is twice as slow as the others.  
 317

318 **5. Two-level nonlinear RAS (2L-NLRAS).** In order to improve the con-  
 319 vergence of NLRAS and ensure the scalability of the method, an additional coarse  
 320 correction can be incorporated. Since the problem is nonlinear, we need to use a non-  
 321 linear coarse grid correction. We opt here for the FAS (Full Approximation Scheme)  
 322 method to perform the coarse grid correction; see, e.g., [8]. Let  $V_0$  denote the coarse  
 323 space and  $R_0: V \mapsto V_0$  the coarse restriction operator. We can define then by means

324 of  $R_0$  the coarse nonlinear problem  $F_0: V_0 \mapsto V_0$  which could be defined by using a  
 325 coarse discretization of (2.1), or using a Galerkin approach, namely

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

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

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

328 which can be incorporated into (2.7) either in an additive or multiplicative manner.  
 329 Following [23], here we opt to use the additive variant. In order to prevent overcor-  
 330 rection on the overlap, we use a weighted two-level method of the form

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

331 The Full Approximation Scheme (FAS) [7] is a particular case of the nonlinear multi-  
 332 level method (NMLM) of Hackbusch [25]. Although NMLM is a globally convergent  
 333 iteration (cf. [25]), there exist no such result for FAS. However, there exist cases for  
 334 which the convergence of FAS was proven for a class of mildly nonlinear PDEs [29,  
 335 Theorem 5.1].

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

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

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

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

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

343 Moreover, we know from Section 4 that

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

344 It remains to compute the derivative corresponding to the coarse correction  $v_0$ . In-  
 345 deed, taking the derivative on both sides of (5.1) with respect to  $u$ , we obtain

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

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

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

348 In order to prove convergence of this two-level method, it suffices then to show that  
 349  $\tilde{T}'(u^*)$  as in (5.5) is a contraction. We can use the same techniques as for the one-level  
 350 method in Theorem 2.2. One added complication is that unlike, e.g., in [23], for our  
 351 coarse grid restriction  $R_0^\top R_0 \neq I$ , and thus additional hypotheses are needed. Then,  
 352 we can use arguments similar to those in the proof of [4, Theorem 7.3] to complete  
 353 the proof, as we show below.

354 **THEOREM 5.1.** *Let  $u^*$  be the solution of (2.2). Assume that  $F'$  exists on a*  
 355 *neighborhood of  $u^*$  and that  $F'(u^*)$  and  $F'(R_0^\top R_0 u^*)$  are nonsingular  $M$ -matrices.*  
 356 *Let  $w > 0$  be such that  $F'(u^*)w > 0$ . Then  $\rho(\tilde{T}'(u^*)) < \|\tilde{T}'(u^*)\|_w < 1$ , and the*  
 357 *iteration (5.2) converges linearly to  $u^*$ .*

358 *Proof.* Let us re-write the local contraction operator (5.5) at the root  $T'(u^*)$  as

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

359 where  $B = \sum_{i=1}^P R_{i,0}^\top (R_{i,\delta} F'(u^*) R_{i,\delta}^\top)^{-1} R_{i,\delta}$ , and  $B_0 = R_0^\top (R_0 F'(R_0^\top R_0 u^*) R_0^\top)^{-1} R_0$ .  
 360 Since  $B$  and  $B_0$  are restrictions and prolongations of nonnegative matrices, where  
 361 the restrictions and prolongation operators are nonnegative, they are nonnegative.  
 362 Furthermore,  $B \geq 0$  is nonsingular, in fact, it is a RAS preconditioner. Therefore,  
 363 it cannot have a zero row, and thus  $Be > 0$ , for any positive vector  $e$ . Let  $e =$   
 364  $F'(u^*)w > 0$ . Then, we have that  $Be > 0$  and  $B_0e \geq 0$ , so that

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


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

366 We present now several numerical experiments comparing the performance of two-  
 367 level NLRAS with the one-level counterpart for problem (4.3), and conduct scaling  
 368 experiments by varying the number of subdomains.

369 We show in Figure 5.1 a weak scaling experiment. We compare the one- and two-  
 370 level NLRAS for a  $2 \times 2$  and  $3 \times 3$  decomposition. The total number of DOFs in each  
 371 subdomain is kept the same for both cases, namely 625 variables per subdomain. As  
 372 we can observe, the convergence of the one-level NLRAS deteriorates when the number  
 373 of subdomains increases from  $P = 4$  to  $P = 9$ . In contrast, the convergence rate of the  
 374 two-level NLRAS remains approximately the same. We summarize weak scalability  
 375 results with increasing numbers of subdomains for iteration (5.2) in Table 5.1, where  
 376 we report number of iterations to reach the tolerance of  $10^{-8}$ . These results show  
 377 that the method without coarse correction fail to converge in 100 iterations, while for  
 378 the two-level NLRAS, the number of iterations has a slow growth with the number of  
 379 subdomains.

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

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



**Fig. 5.1:** Weak scaling experiment for one- and two-level NLRAS for problem (5.6) 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   | 2500  | 36      | >100  | 62       |
| 9   | 5625  | 81      | >100  | 64       |
| 16  | 10000 | 144     | >100  | 94       |
| 25  | 15625 | 225     | >100  | 85       |

**Table 5.1:** Number of iterations performed by NLRAS and two-level NLRAS for the problem (4.3) for different numbers of subdomains  $P$  and a tolerance of  $\epsilon = 10^{-8}$

382 We summarize in Table 5.2 the convergence results of iteration (5.2) for this problem.  
 383 We point out that for a discretization of (5.6) using finite differences,  $F'(u)$  is a  
 384 nonsingular  $M$ -matrix for  $u \in \Omega$ , and therefore the hypothesis of our local convergence  
 385 theorems are satisfied in this case. The results in Table 5.2 are similar to those in  
 386 Table 5.1. The addition of the coarse correction makes the algorithm scalable in the  
 387 sense that the number of iterations required to achieve the prescribed tolerance has  
 388 a very slow growth as the number of subdomains increases. Comparing Tables 5.1  
 389 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.

| $P$ | $ V $ | $ V_0 $ | NLRAS | 2L-NLRAS |
|-----|-------|---------|-------|----------|
| 4   | 2500  | 36      | >100  | 63       |
| 9   | 5625  | 81      | >100  | 64       |
| 16  | 10000 | 144     | >100  | 97       |
| 25  | 15625 | 225     | >100  | 89       |

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

390  
 391 We end this section with two practical observations on the implementation of  
 392 this two-level NLRAS method. The coarse correction step in (5.1) requires the use of



393 Newton's iteration, which in turn necessitates the knowledge of the action of  $F(u)$  as  
 394 well as  $F'(u)$ . These quantities can be obtained by means of the local functions  $F_{i,\delta}$   
 395 as follows. By recalling the definition of  $F_{i,\delta}$ , we have that

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

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

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

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

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

402 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  
 403 nonlinear residual  $R_{i,\delta} F(u^k)$  is zero outside the overlap. This shows that a good coarse  
 404 grid should use this information for the construction of the coarse space. Thus, the  
 405 coarse grid should be chosen on the overlap.

406 **6. Asynchronous Two-Level NLRAS.** In this section, we study the asyn-  
 407 chronous version of the iteration (5.2). We present implementation details, local  
 408 convergence theory, and numerical experiments. The information for the coarse grid  
 409 correction comes from all subdomains, but it cannot be used until the information  
 410 from all subdomains has arrived. Only then, the correction is added, and the process  
 411 repeats. We implement this following ideas described in [23] for linear problems. We  
 412 use a weighted additive NLRAS method, and in order to prevent the coarse correc-  
 413 tion from over-correcting, we add the coarse grid only when all the subdomains have  
 414 sent new data. This can be realized by defining a local boolean variable for each  
 415 thread/subdomain. The coarse grid is then added when all these local boolean vari-  
 416 ables are set to true. We describe the asynchronous two-level NLRAS in Algorithm 2.

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

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

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

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

---

**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:**  $u^*$ .


---

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

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


426 otherwise. Of course the operators (6.1) are those of the two-level NLRAS method (5.3),  
 427 while those in (6.2) are those of the one-level method (2.8). It follows that both opera-  
 428 tors have  $u^*$  as a common fixed point, and in both cases,  $\|T^k\|_w < 1$ , by Theorems 2.2  
 429 and 5.1. Thus, from Theorem 3.3, Algorithm 2 converges to  $u^*$ .  $\square$

430 We compare numerically the parallel synchronous and asynchronous two-level  
 431 NLRAS for the nonlinear problem (4.3). We keep the same setting used in Section 4.  
 432 In particular, the Jacobians are computed exactly for the inner Newton's solves. The  
 433 coarse grid is chosen using a full weighting restriction matrix. We use 10000 dis-  
 434 cretization points for the domain  $\Omega = (0, 1)^2$  which we partition into 16 subdomains.  
 435 The total size of the coarse grid is 144 discretization points. We test our methods in  
 436 two settings: first with all processors of the same speed, and then, with one processor  
 437 slower than the others, in this case, four times slower. These experiments are reported



**Fig. 6.1:** Convergence curves of synchronous and asynchronous two-level NLRAS for problem (4.3) with  $P = 16$ . Left: Homogeneous network. Right: One processor four times slower than the other ones .

438 in Figure 6.1. We can clearly see that in the case of a inhomogeneous network the asynchronous two-level method is faster than its synchronous counterpart. For our



**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 (4.3). Right: Problem (5.6).

439  
 440 weak scaling study, varying the number of subdomains, we report in Figure 6.2  
 441 the time required to reach a tolerance of  $\epsilon = 10^{-8}$  for the synchronous and asynchronous  
 442 two-level NLRAS for problems (4.3) and (5.6). We keep the computation unbalanced  
 443 where one processor is four times slower than the others. We can clearly see that the  
 444 asynchronous method converges faster than the synchronous one.

445 We conclude with a table comparing execution times and the number of processor  
446 updates for the one-level and two-level methods in both the synchronous and asyn-  
447 chronous versions for problem (4.3) in 2D and 3D using  $P + 1$  processors. In the 3D  
448 case, the domain  $\Omega = (0, 1)^3$  is decomposed into  $P$  equally sized bricks, and their  
449 number is the same in each direction. The size of each local problem is  $1000 \times 1000$ .  
450 We assume that one processor is four times slower than the others. The convergence  
451 results are reported in Table 6.1. We mention that we used a smaller coarse grid for  
452 the 64 subdomains to prevent having a coarse grid larger than the size of the local  
453 subdomains. We see from Table 6.1 that the asynchronous method outperforms the  
454 synchronous NLRAS. The addition of the coarse grid correction makes both methods  
455 even faster, with an advantage for the asynchronous one. Observe that as expected,  
456 the asynchronous method has a larger number of updates on average than the number  
457 of iterations of the synchronous method. Nevertheless, its execution time is faster,  
458 confirming that it can outperform the synchronous method. This is true for both the  
one- and two-level methods.

| P  | NLRAS        |       |               |                     | 2L-NLRAS     |       |               |                     |
|----|--------------|-------|---------------|---------------------|--------------|-------|---------------|---------------------|
|    | sync<br>time | #iter | async<br>time | average<br>#updates | sync<br>time | #iter | async<br>time | average<br>#updates |
| 4  | 4.44         | 160   | 3.38          | 374                 | 1.71         | 62    | 1.39          | 131                 |
| 9  | 9.67         | 323   | 5.22          | 581                 | 1.87         | 64    | 1.40          | 153                 |
| 16 | 16.42        | 554   | 7.18          | 880                 | 2.86         | 94    | 1.47          | 170                 |
| 25 | 26.15        | 853   | 10.98         | 1339                | 2.82         | 85    | 1.68          | 187                 |
| 8  | 21.53        | 72    | 17.28         | 163                 | 17.64        | 56    | 14.17         | 150                 |
| 27 | 47.51        | 147   | 30.71         | 268                 | 24.23        | 77    | 17.71         | 159                 |
| 64 | 69.66        | 251   | 58.48         | 468                 | 24.04        | 85    | 22.45         | 178                 |

**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 (4.3) in 2D (top) and 3D (bottom) to reach a tolerance  $\epsilon = 10^{-8}$  for different numbers of subdomains with constant local problem size.

459

460 **7. Conclusion.** In summary, the new parallel nonlinear RAS algorithm con-  
461 verges to the unique solution of the problem in the original domain of the PDE. An  
462 asynchronous version can be about more than twice as fast, depending on the inho-  
463 mogeneities of the computation, and the communication costs. The two-level version  
464 can be an order of magnitude faster, and its asynchronous version faster still, exhibit-  
465 ing slow growth in execution times when the number of processors increases in weak  
466 scaling.

467

## REFERENCES

- 468 [1] Jacques Mohcine Bahi, Sylvain Contassot-Vivier, and Raphaël Couturier. *Parallel Iterative*  
469 *Algorithms: from Sequential to Grid Computing*. CRC Press, 2007.  
470 [2] Zhong-Zhi Bai, Violeta Migallón, José Penadés, and Daniel B. Szyld. Block and asynchronous  
471 two-stage methods for mildly nonlinear systems. *Numerische Mathematik*, 82:1–20, 1999.  
472 [3] Gérard M. Baudet. Asynchronous iterative methods for multiprocessors. *Journal of the Asso-*  
473 *ciation for Computing Machinery*, 25:226–244, 1978.

- 474 [4] Michele Benzi, Andreas Frommer, Reinhard Nabben, and Daniel B. Szyld. Algebraic theory of  
475 multiplicative Schwarz methods. *Numerische Mathematik*, 89(4):605–639, 2001.
- 476 [5] Dimitri P. Bertsekas and John N. Tsitsiklis. *Parallel and distributed computation: numerical*  
477 *methods*, volume 23. Prentice hall Englewood Cliffs, NJ, 1989.
- 478 [6] Amit Bhaya, Eugenius Kaszurewicz, and Francisco Mota. Asynchronous block-iterative meth-  
479 ods for almost linear equations. *Linear Algebra and its Applications*, 154:487–508, 1991.
- 480 [7] Achi Brandt. Multi-level adaptive solutions to boundary-value problems. *Mathematics of*  
481 *Computation*, 31(138):333–390, 1977.
- 482 [8] Achi Brandt. Guide to multigrid development. In Wolfgang Hackbusch and Ulrich Trottenberg,  
483 editors, *Multigrid methods*, pages 220–312. Springer, 1982.
- 484 [9] Xiao-Chuan Cai and David E. Keyes. Nonlinearly preconditioned inexact Newton algorithms.  
485 *SIAM Journal on Scientific Computing*, 24(1):183–200, 2002.
- 486 [10] Xiao-Chuan Cai and Xuefeng Li. Inexact Newton methods with restricted additive Schwarz  
487 based nonlinear elimination for problems with high local nonlinearity. *SIAM Journal on*  
488 *Scientific Computing*, 33(2):746–762, 2011.
- 489 [11] Xiao-Chuan Cai and Marcus Sarkis. A restricted additive Schwarz preconditioner for general  
490 sparse linear systems. *SIAM Journal on Scientific Computing*, 21(2):792–797, 1999.
- 491 [12] Ming Chau, Didier El Baz, Ronan Guivarch, and Pierre Spiteri. MPI implementation of parallel  
492 subdomain methods for linear and nonlinear convection–diffusion problems. *Journal of*  
493 *Parallel and Distributed Computing*, 67(5):581–591, 2007.
- 494 [13] Ljubomir B. Ćirić. Generalized contractions and fixed-point theorems. *Publications de l’Institut*  
495 *Mathématique (Beograd)(NS)*, 12(26):19–26, 1971.
- 496 [14] Victorita Dolean, Martin J. Gander, Walid Kheriji, Felix Kwok, and Roland Masson. Nonlinear  
497 preconditioning: How to use a nonlinear Schwarz method to precondition Newton’s method.  
498 *SIAM Journal on Scientific Computing*, 38(6):A3357–A3380, 2016.
- 499 [15] Maksymilian Dryja and Wolfgang Hackbusch. On the nonlinear domain decomposition method.  
500 *BIT Numerical Mathematics*, 37(2):296–311, 1997.
- 501 [16] Evridiki Efstathiou and Martin J. Gander. Why restricted additive Schwarz converges faster  
502 than additive Schwarz. *BIT Numerical Mathematics*, 43(5):945–959, 2003.
- 503 [17] Mouhamed Nabih El Tarazi. Some convergence results for asynchronous algorithms. *Nu-*  
504 *merische Mathematik*, 39(3):325–340, 1982.
- 505 [18] Andreas Frommer and Hartmut Schwandt. Asynchronous parallel methods for enclosing solu-  
506 tions of nonlinear equations. *Journal of Computational and Applied Mathematics*, 60(1-  
507 2):47–62, 1995.
- 508 [19] Andreas Frommer and Daniel B Szyld. On asynchronous iterations. *Journal of Computational*  
509 *and Applied Mathematics*, 123(1-2):201–216, 2000.
- 510 [20] Andreas Frommer and Daniel B. Szyld. An algebraic convergence theory for restricted addi-  
511 tive Schwarz methods using weighted max norms. *SIAM Journal on Numerical Analysis*,  
512 39:463–479, 2001.
- 513 [21] Luc Giraud and Pierre Spiteri. Résolution parallèle de problèmes aux limites non linéaires.  
514 *ESAIM: Mathematical Modelling and Numerical Analysis*, 25(5):579–606, 1991.
- 515 [22] Christian Glusa, Erik G. Boman, Edmond Chow, Sivasankaran Rajamanickam, and Daniel B.  
516 Szyld. Scalable Asynchronous Domain Decomposition Solvers. *SIAM Journal on Scientific*  
517 *Computing*, pages C384–C409, 2020.
- 518 [23] Gaël Guennebaud, Benoit Jacob, et al. Eigen. URL: <http://eigen.tuxfamily.org>, 2010.
- 519 [24] Wolfgang Hackbusch. *Multigrid Methods and Applications*, volume 4 of *Springer Series in*  
520 *Computational Mathematics*. Srpinger, Berlin Heidelberg, 1985.
- 521 [25] Feng-Nan Hwang and Xiao-Chuan Cai. Improving Robustness and Parallel Scalability of  
522 Newton Method Through Nonlinear Preconditioning. In Ralf Kornhuber, Ronald Hoppe,  
523 Jacques Périaux, Olivier Pironneau, Olof Widlund, and Jinchao Xu, editors, *Domain De-*  
524 *composition Methods in Science and Engineering*, volume 40 of *Lecture Notes in Compu-*  
525 *tational Science and Engineering*, pages 201–208, Berlin, Heidelberg, 2005. Springer.
- 526 [26] Olga A. Ladyzhenskaya and Nina N. Ural’tseva. *Linear and Quasilinear Elliptic Equations*.  
527 *Academic Press, New York*, 1968.
- 528 [27] Jean-Claude Miellou, Didier El Baz, and Pierre Spiteri. A new class of asynchronous iterative  
529 algorithms with order intervals. *Mathematics of Computation*, 67(221):237–255, 1998.
- 530 [28] Arnold Reusken. Convergence of the multigrid full approximation scheme for a class of elliptic  
531 mildly nonlinear boundary value problems. *Numerische Mathematik*, 52(3):251–277, 1987.
- 532 [29] Pierre Spiteri. Parallel asynchronous algorithms: A survey. *Advances in Engineering Software*,  
533 149:102896, 2020.
- 534 [30] Pierre Spiteri, Jean-Claude Miellou, and Didier El Baz. Parallel asynchronous Schwarz and  
535 multisplitting methods for a nonlinear diffusion problem. *Numerical Algorithms*, 33(1-

- 536 4):461–474, 2003.
- 537 [31] Daniel B. Szyld. Different models of parallel asynchronous iterations with overlapping blocks.  
538 *Computational and Applied Mathematics*, 17:101–115, 1998.
- 539 [32] Jian-Jun Xu. Convergence of partially asynchronous block quasi-Newton methods for nonlinear  
540 systems of equations. *Journal of Computational and Applied Mathematics*, 103(2):307–321,  
541 1999.