

**A Two-Level Restricted Additive Schwarz
Method For Asynchronous Computations**

Fayçal Chaouqui and Daniel B. Szyld

Report 22-11-11

Department of Mathematics
Temple University
Philadelphia, PA 19122

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

A Two-Level Restricted Additive Schwarz Method For Asynchronous Computations

Fayçal Chaouqui and Daniel B. Szyld

Abstract The focus of this paper is to compare parallel synchronous and asynchronous domain decomposition solvers for PDEs. An important component that we will address is a coarse space correction for asynchronous computations. We need to ensure that the coarse grid allows the method to converge and work as a solver. In this paper, we propose to use a very specific coarse grid for the asynchronous Restricted Additive Schwarz method. In particular, we compare experimentally both the synchronous and asynchronous methods using this coarse grid.

1 Introduction

In this paper, we study the parallel performance of both synchronous and asynchronous domain decomposition methods for the solution of elliptic PDEs of the form $L(u) = f$ on $\Omega \subset \mathbb{R}^2$ such that $u|_{\Omega} = 0$. The operator L after discretization yields a large sparse system of algebraic equations of the form

$$A\mathbf{u} = \mathbf{f}, \tag{1}$$

where $A \in \mathbb{R}^{n \times n}$ and $\mathbf{f} \in \mathbb{R}^n$. Here, we focus our attention on the Restricted Additive Schwarz (RAS) domain decomposition solver [2, 4]. For the sake of simplicity, we assume that $L = -\Delta$. We assume that the domain Ω is decomposed into p overlapping subdomains $\Omega_1, \dots, \Omega_p$. Let R_i^\top denotes the boolean matrix that maps the local degrees of freedom defined in Ω_i to Ω . We define the local stiffness matrix $A_i = R_i A R_i^\top$. Let us also define the diagonal matrices $D_i, i = 1, \dots, p$ such that we satisfy the partition of unity, i.e.,

Fayçal Chaouqui
COMSOL, INC, Burlington, Mass., USA, e-mail: chaouqui@temple.edu

Daniel Szyld
Temple University, Philadelphia, Pa., USA, e-mail: szyld@temple.edu

$$\sum_{i=1}^p R_i^\top D_i R_i = I, \quad (2)$$

where I denotes the identity matrix in $\mathbb{R}^{n \times n}$. The RAS iteration is then defined as

$$\mathbf{u}^{k+1} = \mathbf{u}^k + \sum_{i=1}^p R_i^\top D_i A_i^{-1} R_i (\mathbf{f} - A \mathbf{u}^k). \quad (3)$$

We note that in our case, the matrices D_i correspond to diagonal boolean matrices that are 1 in the non-overlapping partition, and 0 otherwise. We note also that there are other ways for choosing those matrices, and we refer the reader to [4, 5]. In the next section, we will describe briefly the asynchronous RAS method.

2 Asynchronous Restricted additive Schwarz

Asynchronous methods are parallel iterative procedures where each process performs its task without waiting for other processes to be completed, i.e., with whatever information is available and with no synchronizations with other processes. In this manner, there is no wait, and the communication time is concurrent with computational times. Therefore, there is no communication bottleneck; see, e.g., [1, 3].

We briefly describe asynchronous iterations for fixed point problems defined on a product space $U = U_1 \times \dots \times U_p$, of the form $\mathbf{u} = \mathcal{T} \mathbf{u}$, with a unique solution. In other words, we have $\mathbf{u} = (\mathbf{u}_1, \dots, \mathbf{u}_p)$ and $\mathcal{T} = (\mathcal{T}_1, \mathcal{T}_2, \dots, \mathcal{T}_p)$, with $\mathcal{T}_s : U \rightarrow U_s$. We have in mind that the operation in process s , of the form $\mathbf{u}_s = \mathcal{T}_s(\mathbf{u}_1, \dots, \mathbf{u}_p)$ is performed without synchronization, i.e., without waiting for other processors to send new information.

For a mathematical model of these asynchronous iterations on p processors, we follow the model introduced by Bertsekas [1, 9]. To that end, we define a time stamp k , $k \in \mathbb{N}$, and denote by $\{\sigma(k)\}_{k \in \mathbb{N}}$ the sequence of non-empty subsets of $\{1, \dots, p\}$, defining which processes update their components at the time stamp k . Define also for $s, q \in \{1, \dots, p\}$, $\{\tau_q^s(k)\}_{k \in \mathbb{N}}$ a sequence of integers, representing the update number (or time stamp) of the data coming from process q and available on process s at the time k . Thus, a *delay* would be $k - \tau_q^s(k)$. We begin with an initial approximation $\mathbf{u}^0 = (\mathbf{u}_1^0, \dots, \mathbf{u}_p^0)$, and define, for each process s , the asynchronous iterations as follows.

$$\mathbf{u}_s^{k+1} = \begin{cases} \mathcal{T}_s \left(\mathbf{u}_1^{\tau_1^s(k)}, \dots, \mathbf{u}_p^{\tau_p^s(k)} \right) & \text{if } s \in \sigma(k+1) \\ \mathbf{u}_s^k & \text{if } s \notin \sigma(k+1). \end{cases} \quad (4)$$

In this model, one also assumes that the three following natural conditions are satisfied

$$\forall s, q \in \{1, \dots, p\}, \forall k \in \mathbb{N}, \tau_q^s(k) \leq k, \quad (5)$$

$$\forall s \in \{1, \dots, p\}, \text{card} \{k \in \mathbb{N} | s \in \sigma(k)\} = +\infty, \quad (6)$$

$$\forall s, q \in \{1, \dots, p\}, \lim_{k \rightarrow +\infty} \tau_q^s(k) = +\infty. \quad (7)$$

Condition (5) represents the fact that data used at the time k must have been produced before time k , i.e., time does not flow backward. Condition (6) indicates that no process will ever stop updating its components. Condition (7) means that new data will always be provided to the process. In other words, no process will have a piece of data that is never updated.

One important theoretical result states that for a fixed point problem, say $T(\mathbf{u}) = \mathbf{u}$, on a product space, under conditions 5–7, if there is a weighted max norm such that the map T is contracting, i.e., if the (synchronous) fixed point iteration converges, then, the corresponding asynchronous iteration converges as well; see, e.g., [3] and references therein. For the RAS iteration, the map \mathcal{T}_s defined in (4) is equivalent to

$$\mathcal{T}_s(\mathbf{u}_1, \mathbf{u}_2, \dots, \mathbf{u}_p) = \mathbf{u}_s + R_s \sum_{i=1}^p R_i^\top D_i A_i^{-1} \mathbf{r}_i, \quad (8)$$

where $\mathbf{r}_i = R_i(\mathbf{f} - \mathbf{A}\mathbf{u})$ is the local residual for the subdomain Ω_i , $i = 1, \dots, p$. The implementation of iteration (4) is presented in Algorithm 1.

Algorithm 1 (Asynchronous RAS)

```

1: Input:  $\mathbf{u}^0$ .
2: Output:  $\mathbf{u} \approx \mathbf{u}^*$ .
3: Set  $\mathbf{r}^0 = \mathbf{f} - \mathbf{A}\mathbf{u}^0$ , converged = false.
4: In parallel, each processor  $s$ :
5: while converged = false do
6:   Set  $\mathbf{u}_s = \mathcal{T}_s(\mathbf{u}_1, \dots, \mathbf{u}_p)$  ▷ Update subdomain  $s$ 
7:   Compute  $\|D_s \mathbf{r}_s\|_2$  ▷ Compute local residual norm
8:   if  $s = 1$  then
9:     Compute  $\|\mathbf{r}\|_2 = \sqrt{\sum_{i=1}^p \|D_i \mathbf{r}_i\|_2^2}$ 
10:    if  $\|\mathbf{r}\|_2 / \|\mathbf{r}^0\|_2 \leq \epsilon$  then ▷ Check global convergence
11:      converged = true
12:    end if
13:  end if
14: end while (for processor  $s$ )
15: Set  $\mathbf{u} = \sum_{s=1}^p R_s^\top D_s \mathbf{u}_s$  ▷ Assemble global solution

```

In Algorithm 1 each processor computes and updates the components of the local vector as well as the corresponding local residual norms. A processor is then in charge of accumulating all the local residuals and computing the global residual. The algorithm then stops when the global residual is smaller than the tolerance. In this part, we provide the results of numerical examples illustrating the performance of both synchronous and asynchronous RAS. We consider $\Omega = [0, 1] \times [0, 1]$ decomposed into a checkerboard partition with a total of p subdomains and a minimal

overlap. The source term \mathbf{f} is chosen such that $\sin(\pi x) \sin(\pi y)$ corresponds to the exact monodomain solution. We partition the domain into $p = 16$ subdomains with a total of 10k discretization points. We note that each processor was assigned to one subdomain. All the tests were carried out on a shared memory machine which consists of 88 CPU cores / 176 threads and 1536GB of RAM. The implementation of Algorithm 1 was in C++ and the parallelization uses the OpenMP multithreading directives. We run two different types of experiments. In the first run, we assume all processors run at the same speed and compare both the timings required by both synchronous and asynchronous to reach a specified tolerance. This is illustrated in Figure 2 (left). We can see that in this case the synchronous is faster than the asynchronous. To show the advantage of the asynchronous approach, we repeat the experiment but with one thread slower than the others. This can be experiment can be realized by forcing one processor to be idle for a certain amount of time. We can observe from Figure 2 (right) that the asynchronous is now faster than the synchronous.

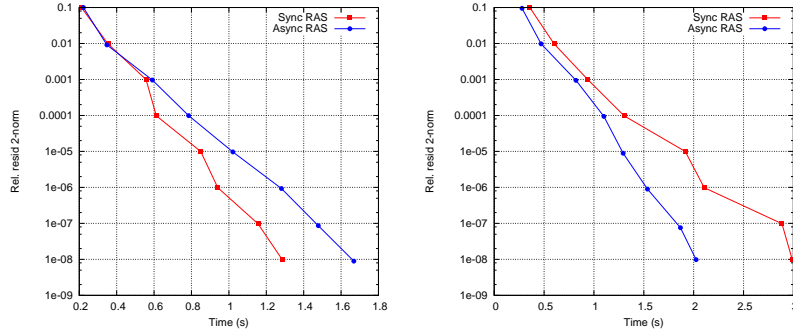


Fig. 1 Left: CPU time versus relative residual 2-norm for synchronous and asynchronous RAS with $p = 16$. Right: Same but one thread is twice as slow.

3 Two-Level Asynchronous Restricted Additive Schwarz

A second level is an essential component to obtain a robust domain decomposition method. It relies generally on solving a smaller problem on a coarser mesh so that there is global communication between the subdomains. The coarse space allows us then to construct the coarse restriction matrix R_0 . The two-level RAS is then defined as

$$\begin{aligned} \mathbf{u}^{k+1/2} &= \mathbf{u}^k + \sum_{i=1}^p R_i^\top D_i A_i^{-1} R_i (\mathbf{f} - A \mathbf{u}^k) \\ \mathbf{u}^{k+1} &= \mathbf{u}^{k+1/2} + R_0^\top A_0^{-1} R_0 (\mathbf{f} - A \mathbf{u}^{k+1/2}), \end{aligned} \quad (9)$$

In order to use iteration (9) asynchronously, we need to use the coarse grid in an additive way. This can be done by using a weighted additive version or multiplicative/additive variant of (9). The corresponding two-level mapping \tilde{T} can be expressed in the case of the additive variant as

$$\tilde{T}_s(\mathbf{u}_1, \mathbf{u}_2, \dots, \mathbf{u}_p) = \mathbf{u}_s + R_s \left(\frac{1}{2} \sum_{i=1}^p R_i^\top D_i A_i^{-1} \mathbf{r}_i + \frac{1}{2} R_0^\top A_0^{-1} R_0 \sum_{i=1}^p R_i^\top D_i \mathbf{r}_i \right), \quad (10)$$

or for the multiplicative-additive as

$$\tilde{T}_s(\mathbf{u}_1, \mathbf{u}_2, \dots, \mathbf{u}_p) = \mathbf{u}_s + R_s \left(\sum_{k=1}^p R_k^\top D_k A_k^{-1} \mathbf{r}_k + R_0^\top A_0^{-1} \bar{R}_0 \sum_{i=1}^p R_i^\top D_i \mathbf{r}_i \right), \quad (11)$$

where $\bar{R}_0 = R_0(I - A \sum_{i=1}^p R_i^\top D_i A_i^{-1} R_i)$. We note that the iteration defined by the means of (11) is equivalent to (9). We note also that the action of \bar{R}_0 can be obtained through the local subdomain vectors. In fact,

$$\begin{aligned} \bar{R}_0 \mathbf{r} &= R_0 \sum_{i=1}^p (\mathbf{r}_i - R_i^\top R_i A R_i^\top D_i A_i^{-1} \mathbf{r}_i) \\ &= R_0 \sum_{i=1}^p (\mathbf{r}_i - R_i^\top A_i D_i A_i^{-1} \mathbf{r}_i). \end{aligned} \quad (12)$$

Here we will restrict ourselves to a coarse correction of the form (10). For work using a multiplicative additive variant, we mention [7, 10]. To avoid over-correction from the coarse grid, we have to make sure that no subdomain is corrected again until all the remaining subdomains have updated at least once [8]. We present in Algorithm 2, the implementation of the asynchronous two-level RAS.

We describe now the coarse correction we use. We would like the coarse grid to ensure the scalability of the method as a solver. In the same spirit of [6], we use harmonically extended coarse basis functions. Let n_i denotes the number of cross points for each subdomain $\Omega_i, i = 1, \dots, p$. Let $\varphi_i^j, j = 1, \dots, n_i$ define a piecewise linear function on $\partial\Omega_i$ that is 1 at one cross point and 0 on the others. We define the coarse basis functions $\phi_i^j, j = 1, \dots, n_i, i = 1, \dots, p$ as the solution of

$$\begin{cases} L|_{\Omega_i}(\phi_i^j) = 0, \text{ on } \Omega_i \\ \phi_i^j = \varphi_i^j, \text{ on } \partial\Omega_i. \end{cases} \quad (13)$$

We define our coarse space $\mathcal{Z} \subset \mathbb{R}^d$ as the span of extended coarse functions ϕ_i^j , i.e.,

$$\mathcal{Z} = \text{span} \left\{ R_i^\top \phi_i^j, j = 1, \dots, n_i, i = 1, \dots, p \right\}. \quad (14)$$

The columns of the matrix R_0^\top correspond to a basis of \mathcal{Z} . We show in Table 1 the number of iterations needed to reach a tolerance $\epsilon = 10^{-8}$ with this specific coarse

Algorithm 2 (Asynchronous two-level RAS)

```

1: Input:  $\mathbf{u}^0$ .
2: Output:  $\mathbf{u} \approx \mathbf{u}^*$ .
3: Set  $\mathbf{r}^0 = \mathbf{f} - A\mathbf{u}^0$ , converged = false.
4: Set update[s]=false, and correction[s]=false,  $s = 1, \dots, p$ .
5: In parallel, each processor  $s$ :
6: while converged = false do
7:   if  $s > 0$  then
8:     if correction[s] then                                     ▷ Check if coarse correction is needed
9:       Set  $\mathbf{u}_s = \mathcal{T}_s(\mathbf{u}_1, \dots, \mathbf{u}_p)$                                ▷ Update subdomain  $s$ 
10:      Set correction[s]=false
11:     else
12:       Set  $\mathbf{u}_s = \mathcal{T}_s(\mathbf{u}_1, \dots, \mathbf{u}_p)$                                ▷ Update subdomain  $s$ 
13:       Set update[s]=true
14:     end if
15:     Compute  $\|D_s \mathbf{r}_s\|_2$                                              ▷ Compute local residual norm
16:     if  $s == 1$  then
17:       Compute  $\|\mathbf{r}\|_2 = \sqrt{\sum_{i=1}^p \|D_i \mathbf{r}_i\|_2^2}$ 
18:       if  $\|\mathbf{r}\|_2 / \|\mathbf{r}^0\|_2 \leq \epsilon$  then                               ▷ Check global convergence
19:         converged = true
20:       end if
21:     end if
22:   else
23:     if update[q],  $\forall q = 1, \dots, p$  then                               ▷ Check if all subdomains updated
24:       Compute the coarse correction.
25:       Set correction[i]=true,  $i = 1, \dots, p$ 
26:       Set update[i]=false,  $i = 1, \dots, p$ 
27:     end if
28:   end if
29: end while (for processor  $s$ )
30: Set  $\mathbf{u} = \sum_{s=1}^p R_s^\top D_s \mathbf{u}_s$                                        ▷ Assemble global solution

```

space for (synchronous) RAS as a solver. We can see that the two-level method outperforms the one-level method and is also scalable, i.e., the number of iterations does not increase when we grow the number of subdomains.

Table 1 Weak scalability of two-level RAS

p	n	$\dim \mathcal{Z}$	#iterations (RAS)	#iterations (additive two-level RAS)
16	10000	36	524	83
25	15625	64	793	85
36	22500	100	1117	86
49	30625	144	1497	87
64	40000	196	1931	87

Next, we test the performance of the two-level asynchronous Algorithm 2 by comparing the time needed to reach a specified tolerance. In Figure 2, we plot the timing versus the residual norm for both synchronous and asynchronous two-level methods. We can observe that in this case the synchronous is faster than the asynchronous. We also see that the timing required to converge is faster than for

the one-level method. The introduction of heterogeneity among processors yields a faster asynchronous two-level method. We note that as is the case for the local subdomains, the coarse problem was solved exactly since it is small for the coarse space defined (14).

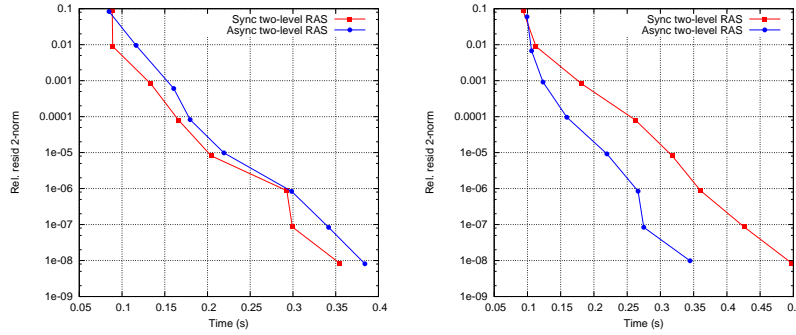


Fig. 2 Left: CPU time versus relative residual 2-norm for two-level synchronous and asynchronous RAS with $p = 16$. Right: Same but one thread is twice as slow.

In Figure 3, we test the weak scalability of both the synchronous and asynchronous methods. To do so, we fix the tolerance to $\epsilon = 10^{-8}$ and run the two-level algorithms and measure the CPU time required to converge. We also plot the time required for the synchronous one as well. In Figure 3 (left), all the processors run at the same speed and there is no load imbalance. We can observe that in this case, the two-level asynchronous method is the fastest among all the four methods. The one-level synchronous method is still slightly faster than the one-level asynchronous. In Figure 3 (right), we repeat the same experiment, but with one processor twice as slow. We can see now that the asynchronous method outperforms the synchronous method. This is true for both the one- and two-level methods. Observe also that while the two-level synchronous method is slightly slower in the simulated heterogeneous architecture (for $p = 64$, 2.4 sec vs. 2.14 sec), the asynchronous method is faster (1.18 sec vs. 1.35 sec). The introduction of heterogeneity clearly shows how asynchronous can be effective in practice.

4 Conclusion

In this paper, we analyzed the performance of one and two-level synchronous and asynchronous RAS. In particular, we used a specific coarse grid correction for our asynchronous computations. Our numerical results suggest that the asynchronous methods exhibit good performance. In particular, we observed that for heterogeneous hardware, the asynchronous outperforms the synchronous method. This was valid for both the one and two-level methods.

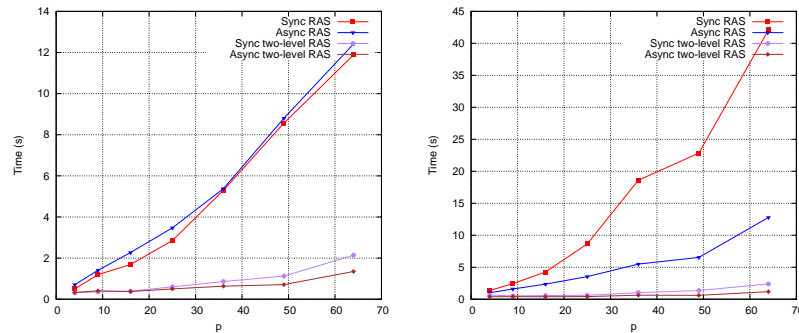


Fig. 3 Left: The number of subdomains versus the CPU time needed for convergence for one and two-level synchronous and asynchronous RAS. Right: Same but with a processor twice as slow.

References

1. Bertsekas, D.P.: Distributed asynchronous computation of fixed points. *Mathematical Programming* **27**(1), 107–120 (1983)
2. Cai, X.C., Sarkis, M.: A restricted additive Schwarz preconditioner for general sparse linear systems. *SIAM Journal on Scientific Computing* **21**, 792–797 (1999)
3. Frommer, A., Szyld, D.B.: On Asynchronous Iterations. *Journal of Computational and Applied Mathematics* **123**, 201–216 (2000)
4. Frommer, A., Szyld, D.B.: An algebraic convergence theory for restricted additive Schwarz methods using weighted max norms. *SIAM Journal on Numerical Analysis* **39**, 463–479 (2001)
5. Gander, M.J.: Does the Partition of Unity Influence the Convergence of Schwarz Methods? In: R. Haynes, S. MacLachlan, X.C. Cai, L. Halpern, H. Kim, A. Klawonn, O. Widlund (eds.) *Domain decomposition methods in science and engineering XXV, Lecture Notes in Computational Science and Engineering*, vol. 138, pp. 3–15. Springer, Cham (2018)
6. Gander, M.J., Loneland, A.: SHEM: An optimal coarse space for RAS and its multiscale approximation. In: C.O. Lee, X.C. Cai, D. Keyes, H. Kim, A. Klawonn, E.J. Park, O. Widlund (eds.) *Domain decomposition methods in science and engineering XXIII, Lecture Notes in Computational Science and Engineering*, vol. 116, pp. 313–321. Springer (2017)
7. Gbikpi-Benissan, G., Magoulès, F.: Asynchronous Multiplicative Coarse-Space Correction. *SIAM Journal on Scientific Computing* **44**, C237–C259 (2022)
8. Glusa, C., Boman, E.G., Chow, E., Rajamanickam, S., Szyld, D.B.: Scalable Asynchronous Domain Decomposition Solvers. *SIAM Journal on Scientific Computing* **42**, C384–C409 (2020)
9. Tsitsiklis, J., Bertsekas, D., Athans, M.: Distributed asynchronous deterministic and stochastic gradient optimization algorithms. *IEEE transactions on automatic control* **31**, 803–812 (1986)
10. Wolfson-Pou, J., Chow, E.: Asynchronous multigrid methods. In: 2019 IEEE international parallel and distributed processing symposium (IPDPS), pp. 101–110. IEEE (2019)