



HAL
open science

Recent Developments in Iterative Methods for Reducing Synchronization

Qinmeng Zou, Frédéric Magoulès

► **To cite this version:**

Qinmeng Zou, Frédéric Magoulès. Recent Developments in Iterative Methods for Reducing Synchronization. 2019 18th International Symposium on Distributed Computing and Applications for Business Engineering and Science (DCABES), Nov 2019, Wuhan, China. pp.163-166, 10.1109/DCABES48411.2019.00048 . hal-02399831v2

HAL Id: hal-02399831

<https://hal.science/hal-02399831v2>

Submitted on 15 Jan 2020

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Recent Developments in Iterative Methods for Reducing Synchronization

Qinmeng Zou*

Frédéric Magoulès*

Abstract. On modern parallel architectures, the cost of synchronization among processors can often dominate the cost of floating-point computation. Several modifications of the existing methods have been proposed in order to keep the communication cost as low as possible. This paper aims at providing a brief overview of recent advances in parallel iterative methods for solving large-scale problems. We refer the reader to the related references for more details on the derivation, implementation, performance, and analysis of these techniques.

Keywords. communication-avoiding algorithms; s -step iterative methods; pipelined Krylov subspace methods; asynchronous iterations.

1 Introduction

The performance of iterative methods depends on the amount of arithmetic operations and the amount of data movements. The latter depends further on the latency cost and the bandwidth cost. Early researches mainly focused on the arithmetic operations in both the sequential and parallel cases. On modern computer architectures, latency cost is much more significant than bandwidth cost, and bandwidth cost is much more significant than computation cost. The gaps are expected to increase in the future.

Parallelization of iterative methods for solving large-scale problems is constrained by synchronization which leads to processor idle time. These methods consist of the following three basic operations: sparse matrix-vector multiplication (SpMV), dot products, and AXPY (α times x plus y) operations

$$y \leftarrow \alpha x + y.$$

AXPY requires only local operations and thus does not affect parallel efficiency. SpMV often requires communication among neighbors, which depends on the distribution of nonzero values. Dot products require global synchronization before and after a computation. The bottleneck that comes from these operations can be partially overcome by using recent techniques. Roughly speaking, the communication in dot products can be reduced by simultaneously constructing multiple direction vectors or being overlapped with other operations; if there is no dot product operation, then the synchronization in SpMV can be eliminated by only using existing data for the next computation instead of waiting for a complete data transmission.

We show in this paper how synchronization points could be reduced and highlight some recent advances on promising techniques. Section 2 presents the s -step iterative methods. Section 3 presents the pipelined Krylov subspace methods. Section 4 summarizes recent developments in asynchronous iterations. Section 5 provides a brief overview of other popular methods. Finally, we draw a conclusion in Section 6.

*CentraleSupélec, Université Paris-Saclay, 3 rue Joliot Curie, 91190 Gif-sur-Yvette, France (zouqinmeng@gmail.com, frederic.magoules@hotmail.com).

2 s -Step Iterative Methods

An early paper describing the idea of s -step iterations can be traced to 1950, as quoted in [26], when Birman [8] presented an s -gradient method in a Russian paper. A later paper [16] was published on the method that aims at reducing the number of synchronization operations for the conjugate gradient (CG) method [35], often called Chronopoulos-Gear CG. The thesis written by Hoemmen [36] gives an excellent historical perspective on this topic. We refer the reader to [36] and the references therein for the developments before 2010.

The key feature in s -step iterative methods is to perform $\mathcal{O}(s)$ computation steps of the classical algorithms for each communication step, thus allowing to reduce the number of synchronization points by a factor of s . Krylov subspace methods (KSMs) are often the methods of choice for solving eigenvalue and linear system problems (see, e.g. [32]), which can be viewed as projection techniques. Given two subspace \mathcal{K}_m and \mathcal{L}_m where

$$\mathcal{K}_m(A, v) = \text{span} \{v, Av, \dots, A^{m-1}v\},$$

KSMs search solution vectors in \mathcal{K}_m such that residuals are orthogonal to \mathcal{L}_m . The latent bottleneck of successive matrix-vector multiplications in KSMs can be relieved by the “matrix powers kernel” as described in [23]. The global synchronization of dot products in Lanczos-based methods can be reduced by using a Gram matrix [11]. In addition, the GMRES method [50] was improved in [36] by using the tall and skinny QR factorization [21].

Hoemmen [36] discussed the s -step GMRES (see also [48]) and s -step CG. Carson et al. [11] discussed the s -step BICG [25] and s -step BICGSTAB [54]. They also addressed the stability issues relating to the basis construction. The term “ s -step” can be often replaced by “communication-avoiding (CA)”. Although the latter is commonly used in the literature, as mentioned in [19], it is slightly dubious since the communication cost is only partially reduced rather than avoided. Ballard et al. [6] summarized theoretical bounds on communication for techniques used in the s -step algorithms.

In finite precision cases, s -step formulation with monomial basis can lead to stability issues, which have been discussed in many references [38, 5, 36, 11]. The maximum attainable accuracy of s -step KSMs and the residual replacement strategy were discussed in [10].

More recently, a new variant proposed by Imberti and Erhel [37] is to use an increasing sequence of block sizes in s -step GMRES instead of a fixed size. On the other hand, block coordinate descent (BCD) methods have been successfully used in machine learning. Devarakonda et al. [24] extended the s -step methods to the primal and dual BCD methods for solving regularized least-square problems. The new methods are called CA-BCD and CA-BDCD, respectively, that can, like other s -step iterative methods, reduce the latency cost by a factor of s but increase computation and bandwidth costs.

3 Pipelined Krylov Subspace Methods

The pipelined iterative methods aim at overlapping expensive communication phases with computations. Some approaches for this purpose applied to CG and GMRES appeared in the mid-1990s, see [22, 20]. Ghysels et al. [30, 31] introduced the modern pipelined Krylov subspace methods. It is interesting to note that the term “pipelined” has often been replaced by “communication-hiding” in their camp, just like the alternatives mentioned in the preceding section.

Ghysels et al. [30] proposed the pipelined GMRES method. Ghysels and Vanroose [31] proposed the pipelined CG method. Their work has promoted other promising ideas. For example, Sanan et al. [51] discussed some pipelined variants of flexible Krylov subspace methods. Cools and Vanroose [18] presented a general framework for pipelined methods, from which the pipelined BICGSTAB method was successfully derived.

In finite precision arithmetic, Cools et al. [19] discussed the effect of local rounding error propagation on the maximal attainable accuracy of the pipelined CG method and compared it with the classical CG and Chronopoulos-Gear CG [16]. In a later paper, Cools [17]

gave a similar discussion for the pipelined BICGSTAB method. Carson et al. [12] discussed the stability issues for synchronization-reducing algorithms and presented a methodology for the theoretical analysis of some CG variants.

4 Asynchronous Iterations

Asynchronous iterations were introduced by Chazan and Miranker [14] in 1969, originally called chaotic iterations. The modern mathematical expression was formally introduced in [7], which can be summarized as follows:

$$x_i^{(n+1)} = \begin{cases} f_i \left(x_1^{(\tau_{i,1,n})}, \dots, x_p^{(\tau_{i,p,n})} \right), & i \in P_n, \\ x_i^{(n)}, & i \notin P_n, \end{cases}$$

where $x_i^{(n)}$ denotes the i th element of the solution vector, $\tau_{i,j,n}$ denotes the iteration number with retards for each element j in each processor i , which is smaller than n , and $P_n \subset \{1, \dots, p\}$ is a subset of processors. In addition, some multi-stage models and a number of convergence results have been proposed during the past century, we refer the reader to [27, 2] for more details. The key feature in asynchronous iterations is to eliminate waiting times in communication at the expense of more iterations.

The main focus of recent research is asynchronous domain decomposition methods. Chau et al. [13] used the asynchronous Schwarz method for the solution of obstacle problems. Magoulès et al. [44] investigated the asynchronous optimized Schwarz method and provided some convergence results. Magoulès and Venet [45] gave a discussion about the asynchronous substructuring methods. More recently, Yamazaki et al. [55] gave some numerical experiments for the optimized Schwarz method. Theoretical analysis of this approach for various types of subdomains were considered in [28, 29]. On the other hand, asynchronous multisplitting methods were studied in [4], which have been applied to the fluid-structure interaction problem in a recent paper [49].

For the time domain decomposition methods, the derivation of asynchronous waveform relaxation can be found in [27]. Magoulès et al. [43, 40] proposed the asynchronous variant of Parareal algorithm (see also [57, 56]). Another asynchronous time-parallel method based on Laplace transform can be found in [46].

From a computational point of view, the implementation of asynchronous iterative methods requires more than a straightforward update of synchronous versions. Magoulès and Gbikpi-Benissan [39, 42] developed an MPI-based communication library for both synchronous and asynchronous iterative computing. However, the issue of asynchronous convergence detection must be tackled for such kind of libraries. Early work can be found in [52] based on the snapshot algorithm. Magoulès and Gbikpi-Benissan [41] continued this work and proposed several promising variants. Bahi et al. [3] introduced another approach in which the detection process is superimposed onto the asynchronous iterations. Numerical experiments of asynchronous iterative methods for GPU were conducted and described in [1, 15].

5 Other Popular Methods

It is clear that other techniques exist for reducing synchronization in parallel architectures, which have not been reviewed in the previous sections. For example, cyclic gradient iterative methods (see, e.g., [58]) are intrinsically adapted for parallel computing, which can reduce both computation and communication costs. Some techniques (see, e.g. [34]) called improved Krylov methods revolve around the overlap of communication and computation. McInnes et al. [47] proposed hierarchical and nested Krylov subspace methods. Grigori et al. [33] proposed enlarged Krylov subspace methods which can be viewed as a special case of augmented Krylov subspace methods (see, e.g, [53]).

6 Conclusion

There is much still to understand about synchronization-reducing methods. For example, development of efficient preconditioners for parallel algorithms is still an open question [9]. Theoretical analysis of basis in s -step algorithms requires more work [36]. The loss of orthogonality in finite precision is also an issue to be tackled [12]. For this reason, we hope that continued contributions could be made on this rapidly growing field in the future.

Acknowledgment

This work was funded by the project ADOM (Méthodes de décomposition de domaine asynchrones) of the French National Research Agency (ANR).

References

- [1] H. Anzt, S. Tomov, J. J. Dongarra, and V. Heuveline. A block-asynchronous relaxation method for graphics processing units. *J. Parallel Distrib. Comput.*, 73(12):1613–1626, 2013.
- [2] J. M. Bahi, S. Contassot-Vivier, and R. Couturier. *Parallel Iterative Algorithms: From Sequential to Grid Computing*. CRC Press, 2007.
- [3] J. M. Bahi, S. Contassot-Vivier, R. Couturier, and F. Vernier. A decentralized convergence detection algorithm for asynchronous parallel iterative algorithms. *IEEE Trans. Parallel Distrib. Syst.*, 16(1):4–13, 2005.
- [4] J. M. Bahi, J.-C. Miellou, and K. Rhofir. Asynchronous multisplitting methods for nonlinear fixed point problems. *Numer. Algorithms*, 15(3):315–345, 1997.
- [5] Z. Bai, D. Hu, and L. Reichel. A Newton basis GMRES implementation. *IMA J. Numer. Anal.*, 14(4):563–581, 1994.
- [6] G. Ballard, E. C. Carson, J. W. Demmel, M. Hoemmen, N. S. Knight, and O. Schwartz. Communication lower bounds and optimal algorithms for numerical linear algebra. *Acta Numer.*, 23:1–155, 2014.
- [7] G. M. Baudet. Asynchronous iterative methods for multiprocessors. *J. ACM*, 25(2):226–244, 1978.
- [8] M. S. Birman. Some estimates for the method of steepest descent. *Uspekhi Mat. Nauk*, 5(3)(37):152–155, 1950. (in Russian).
- [9] E. C. Carson. *Communication-Avoiding Krylov Subspace Methods in Theory and Practice*. PhD thesis, UC Berkeley, 2015.
- [10] E. C. Carson, N. Knight, and J. W. Demmel. A residual replacement strategy for improving the maximum attainable accuracy of s -step Krylov subspace methods. *SIAM J. Matrix Anal. Appl.*, 35(1):22–43, 2014.
- [11] E. C. Carson, N. S. Knight, and J. W. Demmel. Avoiding communication in nonsymmetric Lanczos-based Krylov subspace methods. *SIAM J. Sci. Comput.*, 35(5):S42–S61, 2013.
- [12] E. C. Carson, M. Rozložník, Z. Strakoš, P. Tichý, , and M. Tůma. The numerical stability analysis of pipelined conjugate gradient methods: Historical context and methodology. *SIAM J. Sci. Comput.*, 40(5):A3549–A3580, 2018.

- [13] M. Chau, T. Garcia, and P. Spiteri. Asynchronous Schwarz methods applied to constrained mechanical structures in grid environment. *Adv. Eng. Softw.*, 74:1–15, 2014.
- [14] D. Chazan and W. L. Miranker. Chaotic relaxation. *Linear Algebra Appl.*, 2(2):199–222, 1969.
- [15] E. Chow, H. Anzt, and J. J. Dongarra. Asynchronous iterative algorithm for computing incomplete factorizations on GPUs. In *Proceedings of the 30th International Conference, ISC High Performance 2015*, pages 1–16. Springer, 2015.
- [16] A. T. Chronopoulos and C. W. Gear. s -step iterative methods for symmetric linear systems. *J. Comput. Appl. Math.*, 25(2):153–168, 1989.
- [17] S. Cools. Analyzing and improving maximal attainable accuracy in the communication hiding pipelined BiCGStab method. *Parallel Comput.*, 86:16–35, 2019.
- [18] S. Cools and W. Vanroose. The communication-hiding pipelined BiCGstab method for the parallel solution of large unsymmetric linear systems. *Parallel Comput.*, 65:1–20, 2017.
- [19] S. Cools, E. F. Yetkin, E. Agullo, L. Giraud, and W. Vanroose. Analyzing the effect of local rounding error propagation on the maximal attainable accuracy of the pipelined conjugate gradient method. *SIAM J. Matrix Anal. Appl.*, 39(1):426–450, 2018.
- [20] E. de Sturler and H. A. van der Vorst. Reducing the effect of global communication in GMRES(m) and CG on parallel distributed memory computers. *Appl. Numer. Math.*, 18(4):441–459, 1995.
- [21] J. W. Demmel, L. Grigori, M. Hoemmen, and J. Langou. Communication-optimal parallel and sequential QR and LU factorizations. *SIAM J. Sci. Comput.*, 34(1):A206–A239, 2012.
- [22] J. W. Demmel, M. T. Heath, and H. A. van der Vorst. Parallel numerical linear algebra. *Acta Numer.*, 2:111–197, 1993.
- [23] J. W. Demmel, M. Hoemmen, M. Mohiyuddin, and K. Yelick. Avoiding communication in sparse matrix computations. In *Proceedings of the 22nd International Parallel and Distributed Processing Symposium*, pages 1–12, Miami, FL, USA, 2008. IEEE.
- [24] A. Devarakonda, K. Fountoulakis, J. W. Demmel, and M. W. Mahoney. Avoiding communication in primal and dual block coordinate descent methods. *SIAM J. Sci. Comput.*, 41(1):C1–C27, 2019.
- [25] R. Fletcher. Conjugate gradient methods for indefinite systems. In *Numerical Analysis*, pages 73–89. Springer, 1976.
- [26] G. E. Forsythe. On the asymptotic directions of the s -dimensional optimum gradient method. *Numer. Math.*, 11(1):57–76, 1968.
- [27] A. Frommer and D. B. Szyld. On asynchronous iterations. *J. Comput. Appl. Math.*, 123(1-2):201–216, 2000.
- [28] J. C. Garay, F. Magoulès, and D. B. Szyld. Convergence of asynchronous optimized Schwarz methods in the plane. In *Domain Decomposition Methods in Science and Engineering XXIV*, Lecture Notes in Computer Science and Engineering, pages 333–341. Springer, 2018.
- [29] J. C. Garay, F. Magoulès, and D. B. Szyld. Optimized Schwarz method for Poisson’s equation in rectangular domains. In *Domain Decomposition Methods in Science and Engineering XXIV*, Lecture Notes in Computer Science and Engineering, pages 533–541. Springer, 2018.

- [30] P. Ghysels, T. J. Ashby, K. Meerbergen, and W. Vanroose. Hiding global communication latency in the GMRES algorithm on massively parallel machines. *SIAM J. Sci. Comput.*, 35(1):C48–C71, 2013.
- [31] P. Ghysels and W. Vanroose. Hiding global synchronization latency in the preconditioned conjugate gradient algorithm. *Parallel Comput.*, 40(7):224–238, 2014.
- [32] G. H. Golub and C. F. Van Loan. *Matrix Computations*. Johns Hopkins University Press, 4th edition, 2013.
- [33] L. Grigori, S. Moufawad, and F. Nataf. Enlarged Krylov subspace conjugate gradient methods for reducing communication. *SIAM J. Matrix Anal. Appl.*, 37(2):744–773, 2016.
- [34] T.-X. Gu, X.-Y. Zuo, L.-T. Zhang, W.-Q. Zhang, and Z.-Q. Sheng. An improved bi-conjugate residual algorithm suitable for distributed parallel computing. *Appl. Math. Comput.*, 186(2):1243–1253, 2007.
- [35] M. R. Hestenes and E. Stiefel. Methods of conjugate gradients for solving linear systems. *J. Res. Natl. Bur. Stand.*, 49(6):409–436, 1952.
- [36] M. Hoemmen. *Communication-Avoiding Krylov Subspace Methods*. PhD thesis, UC Berkeley, 2010.
- [37] D. Imberti and J. Erhel. Varying the s in your s -step GMRES. *Electron. Trans. Numer. Anal.*, 47:206–230, 2017.
- [38] W. D. Joubert and G. F. Carey. Parallelizable restarted iterative methods for non-symmetric linear systems. Part I: Theory. *Int. J. Comput. Math.*, 44(1-4):243–267, 1992.
- [39] F. Magoulès and G. Gbikpi-Benissan. JACK: An asynchronous communication kernel library for iterative algorithms. *J. Supercomput.*, 73(8):3468–3487, 2017.
- [40] F. Magoulès and G. Gbikpi-Benissan. Asynchronous Parareal time discretization for partial differential equations. *SIAM J. Sci. Comput.*, 40(6):C704–C725, 2018.
- [41] F. Magoulès and G. Gbikpi-Benissan. Distributed convergence detection based on global residual error under asynchronous iterations. *IEEE Trans. Parallel Distrib. Syst.*, 29(4):819–829, 2018.
- [42] F. Magoulès and G. Gbikpi-Benissan. JACK2: An MPI-based communication library with non-blocking synchronization for asynchronous iterations. *Adv. Eng. Softw.*, 119:116–133, 2018.
- [43] F. Magoulès, G. Gbikpi-Benissan, and Q. Zou. Asynchronous iterations of Parareal algorithm for option pricing models. *Mathematics*, 6(4):1–18, 2018.
- [44] F. Magoulès, D. B. Szyld, and C. Venet. Asynchronous optimized Schwarz methods with and without overlap. *Numer. Math.*, 137(1):199–227, 2017.
- [45] F. Magoulès and C. Venet. Asynchronous iterative sub-structuring methods. *Math. Comput. Simul.*, 145:34–49, 2018.
- [46] F. Magoulès and Q. Zou. Asynchronous time-parallel method based on Laplace transform. *preprint available at arXiv:1909.01473*, 2019.
- [47] L. C. McInnes, B. F. Smith, H. Zhang, and R. T. Mills. Hierarchical Krylov and nested Krylov methods for extreme-scale computing. *Parallel Comput.*, 40(1):17–31, 2014.

- [48] M. Mohiyuddin, M. Hoemmen, J. W. Demmel, and K. Yelick. Minimizing communication in sparse matrix solvers. In *Proceedings of the 22nd Conference on High Performance Computing, Networking, Storage and Analysis*, pages 1–12, Portland, OR, USA, 2009. ACM.
- [49] V. Partimbene, T. Garcia, P. Spiteri, P. Marthon, and L. Ratsifandrihana. Asynchronous multi-splitting method for linear and pseudo-linear problems. *Adv. Eng. Softw.*, 133:76–95, 2019.
- [50] Y. Saad and M. H. Schultz. GMRES: A generalized minimal residual algorithm for solving nonsymmetric linear systems. *SIAM J. Sci. Stat. Comput.*, 7(3):856–869, 1986.
- [51] P. Sanan, S. M. Schnepp, and D. A. May. Pipelined, flexible Krylov subspace methods. *SIAM J. Sci. Comput.*, 38(5):C441–C470, 2016.
- [52] S. A. Savari and D. P. Bertsekas. Finite termination of asynchronous iterative algorithms. *Parallel Comput.*, 22(1):39–56, 1996.
- [53] V. Simoncini and D. B. Szyld. Recent computational developments in Krylov subspace methods for linear systems. *Numer. Linear Algebra Appl.*, 14(1):1–59, 2007.
- [54] H. A. van der Vorst. Bi-CGSTAB: A fast and smoothly converging variant of Bi-CG for the solution of nonsymmetric linear systems. *SIAM J. Sci. Stat. Comput.*, 13(2):631–644, 1992.
- [55] I. Yamazaki, E. Chow, A. Bouteiller, and J. J. Dongarra. Performance of asynchronous optimized Schwarz with one-sided communication. *Parallel Comput.*, 86:66–81, 2019.
- [56] Q. Zou, G. Gbikpi-Benissan, and F. Magoulès. Asynchronous communications library for the parallel-in-time solution of Black-Scholes equation. In *Proceedings of the 16th International Symposium on Distributed Computing and Applications to Business, Engineering and Science*, pages 45–48, Anyang, China, 2017. IEEE.
- [57] Q. Zou, G. Gbikpi-Benissan, and F. Magoulès. Asynchronous Parareal algorithm applied to European option pricing. In *Proceedings of the 16th International Symposium on Distributed Computing and Applications to Business, Engineering and Science*, pages 37–40, Anyang, China, 2017. IEEE.
- [58] Q. Zou and F. Magoulès. Reducing the effect of global synchronization in delayed gradient methods for symmetric linear systems. *Adv. Eng. Softw.*, submitted.