Proceedings of CHT-17
ICHMT International Symposium on Advances in Computational Heat Transfer

May 28-June 1, 2017, Napoli, Italy

CHT-17-021

SPATIAL VERSUS ANGULAR PARALLELIZATION FOR SOLUTION OF RADIATIVE TRANSFER EQUATION IN PARTICIPATING MEDIA

M.A. Badri[*,**], Y. Favennec[**,§], P. Jolivet[***], D. Le Hardy[**], S. Le Corre[**], and B. Rousseau[**]
[*]Institut de Recherche Technologique Jules Verne, 44340 Bouguenais, France
[**]Laboratoire Thermocinétique de Nantes (UMR CNRS 6607), 44306 Nantes, France
[***]Institut de Recherche en Informatique de Toulouse (UMR CNRS 5505), 31062 Toulouse, France
[§] Correspondence author. Ph: +33 (0)240683138 Email: yann.favennec@univ-nantes.fr

ABSTRACT This paper presents two very different methods of parallelization applicable to the solution of the multi-dimensional monochromatic radiative transfer equation. Algorithms concerning the two methods, namely angular decomposition methods and domain decomposition methods, are separately developed using the Message Passing Interface (MPI) and implemented in order to achieve fast and memory efficient solutions. A detailed comparison of performance and scalability on thousands of processors is established for both methods.

INTRODUCTION AND MOTIVATION

Radiative transfer calculations are vital for a wide range of natural and engineering applications, primarily the ones involved in high temperature physics. Many such applications are modelled using the monochromatic steady-state radiative energy, which at mesoscopic and macroscopic scales, is often conceptualized by the transport of radiative intensity $I(\boldsymbol{x}, \boldsymbol{s})$, via an integro-differential equation known as the radiative transfer equation (RTE):

$$(\boldsymbol{s} \cdot \nabla + \kappa + \sigma_s)\, I(\boldsymbol{x}, \boldsymbol{s}) = \sigma_s \oint_{\mathcal{S}^{n-1}} I(\boldsymbol{x}, \boldsymbol{s}')\Phi(\boldsymbol{s}, \boldsymbol{s}')\, \mathrm{d}\boldsymbol{s}' + \kappa I_b(\boldsymbol{x}) \quad \forall \boldsymbol{x} \in \mathbb{R}^n, \boldsymbol{s} \in \mathcal{S}^{n-1} \tag{1}$$

Here $\boldsymbol{x} \in \Omega \subset \mathbb{R}^n$ reads the space variable with $n = 2$ or 3, for 2D or 3D domains, respectively, and $\boldsymbol{s}$ defines the unit vector pointing towards the unit sphere $\mathcal{S}^{n-1}$. On the right-hand side of Eq. (1), κI_b is the Plank's black body emission source, and the source due to scattering effects in radiation is depicted by the integral term. To the left, we sum up loss terms via absorption κI, scattering $\sigma_s I$ and a transport term $\boldsymbol{s} \cdot \nabla I$. Due to its integro-differential nature, RTE cannot be used for computing the radiation field via any explicit numerical method. Hence we follow a two-step discretization process i) angular discretization on angular space $\mathcal{S}^{n-1}$ ($\mathbb{P}_0$ finite element basis functions) and ii) spatial discretization on spatial space Ω ($\mathbb{P}_1$ or $\mathbb{P}_2$ vectorial finite element basis functions). Following the $\mathbb{P}_0$ angular space discretization, the radiative energy $I(\boldsymbol{s}, \boldsymbol{x})$ is discretized into m groups of intensities based on ordinates\angles, which is similar to what we would obtain with the Discrete Ordinate Method (DOM).

Computing the solution of such equations requires solving a large linear system. So far, much work has been performed in the area of parallelization when radiative transfer is solved using stochastic approaches such as Monte Carlo methods. However, the parallelization of deterministic approaches (such as the finite element method) has hardly been studied in this field.

From hereafter the remainder of the paper is organized as follows: a vectorial finite element strategy which is used to reduce the system of equations into a solvable linear system is discussed next. Then a section on parallelization strategies for solving RTE is presented. In the last part of the paper, we discuss and conclude the results obtained via our solution strategies.

VECTORIAL FINITE ELEMENTS FOR SOLVING RTE

When it comes to finite element solutions of RTE, traditionally, the Streamline Upwind Petrov--Galerkin (SUPG--FEM) and the Discontinuous Galerkin--FEM are most widely used. Le Hardy et al. [2016] presents the comparison between the two different methods for a 2D case. These two methods, when applied successively for each direction of the angular discretization, yield a huge block matrix which is composed of $n \times n$ submatrices, each row of the global matrix being related to a particular direction, and the off-diagonal matrices being related to the scattering operator. In such a way, the unknown $I_m(\boldsymbol{x}), \forall m = 1, \ldots, N_d$, is to be searched in a given functional space, say $\mathcal{V}_h$, which is itself a subset of $H^1(\mathcal{D})$.

Alternatively, we introduce SUPG based on vectorial finite elements. Note that this notion, which is also known as the Mixed Finite Element Methods, was first used by Herrmann [1967] in connection with elastic theory to denote methods, based on the Hellinger--Reissner principle, in which both displacements and stresses were approximated simultaneously. The problem consists in searching a vector of radiative intensities $\mathbb{I} = [I_1\ I_2\ \cdots\ I_{N_d}]^T$ utilizing a vectorial test function $\mathbb{V} = [v_1\ v_2\ \cdots\ v_{N_d}]^T$ and, in a similar fashion $\mathbb{S}$ and Θ, for denoting directions and scattering. Doing so, the functional space in which the solution is searched may be, for example, $H^1(\mathcal{D})^{N_d} = \prod_{i=1}^{N_d} H^1(\mathcal{D})$, i.e., a mixed space. Denoting $\beta = (\kappa + \sigma_s)$ as the extinction coefficient, the RTE (1), can be rewritten to its semi-discretized form:

$$\mathbb{S} \cdot \nabla \mathbb{I} + \beta \mathbb{I} - \Theta \mathbb{I} = \kappa I_b \mathbb{1}$$

Then, a SUPG vectorial weak formulation can now be built by multiplying the equation with a vectorial trial function, for example $\mathbb{H} = \mathbb{V} + \gamma \mathbb{S} \cdot \nabla \mathbb{V}$, $\gamma : \Omega \mapsto \mathbb{R}^+$, then integrating it over the domain $\mathcal{D}$

$$\int_{\mathcal{D}} [(\mathbb{S} \cdot \nabla \mathbb{I} + \beta \mathbb{I}) : (\mathbb{H})]\, \mathrm{d}\boldsymbol{x} - \int_{\mathcal{D}} [(\Theta \mathbb{I}) : (\mathbb{H})]\, \mathrm{d}\boldsymbol{x} = \int_{\mathcal{D}} [(\kappa I_b \mathbb{1}) : (\mathbb{H})]\, \mathrm{d}\boldsymbol{x} \tag{2}$$

To follow, Green's theorem is applied on the advection term of Eq. 2, in order to exhibit boundary conditions. A finite element approximation on this equation yields a linear system, $\mathbf{A}\mathbb{I} = \mathbb{b}$, which is then to be solved. Note the large dimension of such matrix, as soon as one considers hundreds of directions.

PARALLEL RTE SOLUTION

Considering the fact that there exists a need to solve large linear system of equations, for time efficiency (during the assembly and the solution phases), parallelization becomes inevitable. Although domain decomposition is the standard approach for numerical parallelization of partial differential equations, the RTE being dependent on both spatial and angular discretizations, we attempt to parallelize it via both domain decomposition (DD) and angular decomposition (AD) approaches.

Spatial Decomposition Algorithm for RTE Contrary to the classical matrix formulated using SUPG--FEM, SUPG vectorial FEM formulates a sparse and banded matrix. We tend to exploit such structure by applying a Krylov subspace method for the solution phase. The original mesh $\mathcal{M}$ is partitioned into N number of small local overlapping meshes $\mathcal{M}_i$. Each process owns one of these local submeshes. A mesh partitioner such as METIS Karypis and Kumar [1995] is used to ensure proper load balancing. The vectorial variational formulation Eq. (2) is used to assemble the local matrices in each MPI process. From here on, we use the Restricted Additive Schwarz (RAS) method as a preconditioner for the GMRES algorithm. Following the notation of Jolivet et al. [2012], we use a set of diagonal matrices $\{D_i\}_{i=1,\ldots,N}$ which defines a partition of unity, and a set of restriction matrices $\{R_i\}_{i=1,\ldots,N}$, such that the preconditioner reads:

$$M^{-1} = \sum_{i=1}^{N} R_i^T D_i (R_i A R_i^T)^{-1} R_i$$

Angular Decomposition Algorithm for RTE In this approach, we dispatch the angular space to different MPI processes. Since the lowest order $\mathbb{P}_0$ finite element discretization in angles leads to discrete equations which can be independently built on the mesh $\mathcal{M}$, this results in an embarrassingly parallel work distribution. Using vectorial finite elements, these equations can be combined in sets leading to sparse matrices for each set of ordinates. AD possess the following advantages: i) perfect load balancing, as all MPI process are in charge of exactly the same number of unknowns; ii) each MPI process independently assembles and solves their respective linear systems without the need of communications; iii) consistent results, as the number of iterations required to achieve convergence is independent of the number of MPI processes. The only limitation of AD comes from the fact that since low-order $\mathbb{P}_0$ elements are used to discretize the angular space (in practice, no more than hundreds of angles), the maximum number of usable MPI process is defined by the discretization.

RESULTS

Two numerical tests involving a 2D geometry with value of scattering albedo $\omega \simeq 0$ (highly absorbing) and $\omega \simeq 1$ (highly scattering), and 32 $\mathbb{P}_0$ unknowns for the angular space discretization, were solved on supercomputer LIGER, at ICI supercomputing facility (6,048 cores, Intel Xeon cluster) in Ecole Central Nantes, France. We assumed a simple square geometry $[0,1] \times [0,1]$ meshed with 21,000 nodes, with Dirichlet boundary conditions in terms of radiative intensities on its left border. It is also worth mentioning that before performing these tests, code arising from both AD and DD algorithms were validated using the method of constructed solutions as used in Le Hardy et al. [2016]. In more details, the codes have been written using the finite element domain-specific library FreeFem++ to assemble the linear systems, which then calls PETSc (Balay et al. [2014]) for solving them.

The aim of the numerical tests is to compare the parallel scaling of AD and DD methods for solving RTE. The comparison is based on elapsed wall-clock time for two major subroutines: matrix assembly and solution phase. We neglect pre- and post-processing times.

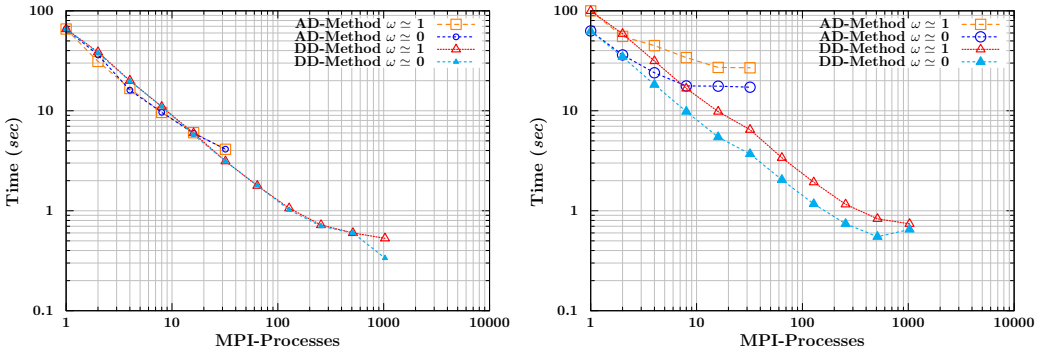

Figure 1. Scaling test comparison for two cases of high and negligible scattering; Left: Matrix building time; Right: GMRES solving time

In Fig. 1, the scalability for both angular and domain decomposition methods can be very well demonstrated: computation time for matrix assembly decreases almost linearly when the number of MPI processes are increased. For solving time, DD methods exhibit superior scalability. Since 32 $\mathbb{P}_0$ unknowns were considered for the angular space discretization, AD algorithm is limited to 32 MPI-processes while there is no such limit for DD method.

From Table 1, notice that, for both AD and DD methods, the total iteration counts remain almost constant while increasing the number of MPI processes. However, total iteration count for DD method remains higher than that of AD method. We don't use any preconditioner yet, but this table suggests that the condition number is rather low since the Krylov solver converges in at most 500 iterations.

Table 1
Comparison between AD and DD method for a problem with 2.62×10^6 degrees of freedom

# of MPI processes	# of iterations ($\omega \simeq 1$)		# of iterations ($\omega \simeq 0$)	
	AD	DD	AD	DD
2	473	516	293	303
4	473	521	293	303
8	473	524	293	305
16	473	523	293	304
32	473	519	293	301
128		518		300
512		518		300
1024		517		300

This is very promising since we should be able to design preconditioners that will make the iterative solver converge very rapidly (in terms of number of iterations).

CONCLUSIONS

We conclude that our approach for both spatial and angular decomposition is scalable on parallel machines. DD method is scalable to higher processor counts while AD has a strong limit, the order of the angular discretizations. Considering no preconditioning for the solving phase, DD methods exhibit superior scalability than AD methods. The design of more efficient preconditioners, which could change the dynamics of time to solution, will be our next objective.

Acknowledgements We would like to thank the staff of the ICI supercomputing facility in Nantes, France, for granting us access to LIGER. M.A. Badri would like to thank IRT Jules Verne (French Institute in Research and Technology in Advance Manufacturing Technologies for Composite, Metallic and Hybrid Structures) for funding his research. He also wishes to associate the industrial and academic partners of the project: Airbus Group Innovations, Dassault Aviation, Safran Composites, GeM (Institut de Recherche en Génie Civil et Mécanique) and LTN (Laboratoire de Thermocinétique de Nantes).

REFERENCES

Balay, S., Abhyankar, S., Adams, M., Brown, J., Brune, P., Buschelman, K., Eijkhout, V., Gropp, W., Kaushik, D., Knepley, M., et al. [2014]. PETSc users manual revision 3.5. *Argonne National Laboratory (ANL)*.

Herrmann, L. R. [1967]. Finite-element bending analysis for plates. *Journal of the Engineering Mechanics Division*, Vol. 93, No. 5, pp 13-26.

Jolivet, P., Dolean, V., Hecht, F., Nataf, F., Prud'Homme, C., and Spillane, N. [2012]. High performance domain decomposition methods on massively parallel architectures with FreeFem++. *Journal of Numerical Mathematics*, Vol. 20, No. 3-4, pp 287-302.

Karypis, G. and Kumar, V. [1995]. METIS--unstructured graph partitioning and sparse matrix ordering system, version 2.0.

Le Hardy, D., Favennec, Y., and Rousseau, B. [2016]. Solution of the 2-D steady-state radiative transfer equation in participating media with specular reflections using SUPG and DG finite elements. *Journal of Quantitative Spectroscopy and Radiative Transfer*, Vol. 179, pp 149-164.