Parallel computation for the radiative heat transfer using the DOM in combustion applications: direction, frequency, sub-domain decompositions and hybrid methods – NHT12/5512

POITOU Damien (poitou@cerfacs.fr)[a], Amaya Jorge[a], Duchaine Florent[a]

[a] *CERFACS, 42, Avenue Gaspard Coriolis, 31057 Toulouse Cedex 01, France.*

Preprint submitted to Numerical Heat Transfer, Part B Fundamentals March 26, 2012

Title Parallel computation for the radiative heat transfer using the DOM in combustion applications: direction, frequency, sub-domain decompositions and hybrid methods – NHT12/5512

Short Title Parallel computation for DOM in combustion

Abstract

In the framework of coupled LES/DOM computations of turbulent combustion problems, different decompositions for parallel calculations of the radiative heat transfer based on the DOM were investigated. The methods analysed are: A) A task decomposition on the discrete directions and frequencies with two numeric strategies: Message Passing Interface (MPI) with distributed memory and a OpenMP with shared memory for the direction decomposition; B) A new algorithm for a DOM sub-domain decomposition is proposed and tested using MPI; C) Hybrid methods combining OpenMP strategy for direction and MPI for tasks and sub-domain decomposition. It was shown in the case of coupled simulations that the convergence and the parallel efficiency of the domain decomposition (B) are optimal. This method is however limited in this work to 25 sub-domains at which point the efficiency stagnates. Combining the directions with frequency and/or domain decompositions in a hybrid method (C) result in a very good efficiency up to 1200 processors. This hybrid strategy is also very efficient in terms of memory usage. This work shows that the best way to perform massively parallel computation for the radiative heat transfer with DOM is to combine different decomposition levels. The analysis performed in this work shows the best parallel strategy to be used in coupled simulations between radiation and LES on massively parallel architectures.

Keywords: Radiative transfer, Discrete ordinate method (DOM), Parallel calculations, Sub-domains, Code coupling, Turbulent combustion, Large Eddy Simulation.

Notations

DMFS	Diamond Mean Flux Scheme
DOM	Discrete Ordinate Method
FSK	Full Spectrum κ (*i.e.* FS-SNBcK)
FSCK	Full Spectrum Correlated κ
FVM	Finite Volume Method
LES	Large Eddy Simulation
RCB	Recursive Bisection Method
RTE	Radiative Transfer Equation
SNB	Statistcal Narrow Band
SNBcK	SNB with correlated κ model
WSGG	Weighted Sum of Gray Gases

1. Introduction

Thermal constraints are critical for the design of industrial combustion chambers and furnaces. Solid walls must be cooled down using complex and expensive techniques to design as materials do not support a direct contact with high temperatures. Thermal behaviour also has an impact on pollutants emission such as CO, NO_x and soot which are very sensitive to the temperature levels. If soot is produced inside the chamber, the transfer of energy is increased due to the radiative contribution of soot. It is therefore necessary to include radiation in combustion studies for the optimization of combustion chambers.

Numerical simulations become more and more important in the aeronautical industry. Important progress have been made in the last few years, in particular with the development of the Large Eddy Simulations (LES) approach and the sustained increase of computational power. Today LES gives accurate and reliable solutions when applied to industrial burners with complex geometries such as combustion helicopter engines, considering a sector of an annular chamber [1] or a full 360° geometry [2]. These kind of calculations can only be done at very high CPU cost and require the use of

High Performance Computing (HPC) architectures. Programming in such an environment remains a challenge even for CFD solvers [3].

Recent works have demonstrated the feasibility of the coupling of a radiative heat transfer solver with LES for unsteady calculations of turbulent combustion [4, 5, 6]. Poitou *et al.* presented the coupling methodology for such applications in detail in [7]. These work has been done using a discrete ordinate method (DOM) solver for the radiative heat transfer with non-gray gases properties on unstructured meshes. The radiative heat transfer depends on non-local exchanges thus the classical sub-domain decomposition used in CFD is not well adapted.

In classical DOM calculations the solution is obtained by an explicit calculation of the radiation intensity field using an upstream propagation of information in the solving direction [8]. This approach is based on a mesh sweeping procedure. Other solvers use a much more simple approach based on an algebraic procedure [9].

Gonçalves and Coelho [10] compared the directions decomposition and the domains decomposition with an algebraic solver. They worked on two-dimensional test cases with a grid size of 90^2 and S_{12} angular quadrature (80 directions) with a maximum of 90 cores, showing a favorable bias towards angular decomposition. Krishnamoorthy *et al.* proposed a review on parallel radiative calculations [9], they also worked on an algebraic solver and a spatial decomposition for different test cases with a maximum grid size of 121^3, showing a good parallel efficiency in gray cases for different optical thickness up to 100 cores. Scaled problems (*i.e.* each processor has the same amount of work to do) were tested with about 1000 cores for non-gray calculations on a more realistic test cases and shown a reduced efficiency, the solve time which should keep constant increases by a about 20 times. In nuclear engineering field, Pautz [11] and Plimpton *et al.* [12] showed the feasibility of sub-domain decomposition with a DOM with an explicit solver based on a parallel sweeping decomposition. This strategy was used recently by Colomer *et al.* in [13] to couple convection and radiation. In these works a parallel decomposition of the sweeping is done, taking into account the dependences between the different nodes by keeping track, for each node, of

the upstream and downstream nodes in each one of the ordinates and all the ordinates are solved simultaneously. Colomer *et al.* showed a good parallel efficiency up to 448 cores on a 64^3 grid with an S_8 angular quadrature (80 ordinates).

In the present work, all these different parallel decomposition methods are compared for an explicit DOM solver in order to identify the best option for the coupled simulation including radiation and LES. The methodology is proposed in the framework of unsteady coupled simulations of turbulent combustion where the radiative calculation is repeated at each coupling points. No coupled calculations will be presented in this paper and the efficiency of the radiative solver is evaluated in a post-processing way. It was shown in [7] that in this context a S_4 angular quadrature (24 ordinates) was enough to get accurate results for aeronautical propulsion applications.

There are usually two ways to decompose a radiative calculation. The first one rely on task decompositions solving the independent parts of a calculation in different cores. The second decomposition method is based on data partitioning. The different decompositions used are shown in Fig. 1 : the directions and the frequencies are used for the task decomposition and sub-domains for the data decomposition.

In this work two numerical strategies are also tested using a Message Passing Interface (MPI) with distributed memory and a shared memory approach using OpenMP. The OpenMP approach was implemented only for the task decomposition method based on the directions, all others decompositions use MPI.

In section 2, the elements of the radiative heat transfer modelling are presented with the numerical implementation and the sweeping ordering. Section 3 presents the two main parallel strategies, tasks decomposition on directions and frequencies and sub-domain decomposition. Section 4 shows the results obtained on a test case of an helicopter combustion chamber. The parallel efficiency is presented for the direction (A1), frequency task decomposition (A2) and the domain decomposition (B). Then the parallel efficiency is investigated for hybrid decomposition using OpenMP for the direction decomposition and MPI for the other ones: direction/frequency (C1),

direction/sub-domains (C2) and direction/frequency/sub-domain (C3). Finally the memory use of the different strategies is given.

2. Modelling of the radiative heat transfer

2.1. Radiative Transfer Equation (RTE)

The RTE is solved in its differential form (Eq. (1)) in the direction of propagation $\underline{\Omega}$, for a non scattering medium, with the associated boundary conditions (Eq (2)):

$$\underline{\Omega} \cdot \nabla L_\nu(\mathbf{x}, \mathbf{u}) = \kappa_\nu \left[L^0_\nu(\mathbf{x}) - L_\nu(\mathbf{x}, \mathbf{u}) \right] \quad (1)$$

$$L_\nu(\mathbf{x}_w, \mathbf{u}) = \underbrace{\epsilon_\nu(\mathbf{x}_w) L^0_\nu(\mathbf{x}_w)}_{\text{Emitted part}} + \underbrace{\rho_\nu(\mathbf{x}_w) L_{\nu,incident}(\mathbf{x}_w, \mathbf{u})}_{\text{Reflected part}} \quad (2)$$

where ν is the wavenumber, $L_\nu(\mathbf{x}, \mathbf{u})$ is the radiation intensity at the point $\mathbf{x}$ in the direction $\mathbf{u}$ and κ_ν is the absorption coefficient, $\epsilon_\nu(\mathbf{x}_w)$ is the wall emissivity and $\rho_\nu(\mathbf{x}_w)$ the wall reflectivity with $\rho_\nu(\mathbf{x}_w) = 1 - \epsilon_\nu(\mathbf{x}_w)$. L^0_ν is the equilibrium Planck function.

The source term S_r is injected in the energy balance equation of the flow and results from a double integration of the RTE over the solid angle and the gas spectra, and depends only on the position $\mathbf{x}$:

$$S_r(\mathbf{x}) = \int_0^\infty \kappa_\nu \left[4\pi L^0_\nu(\mathbf{x}) - \int_{4\pi} L_\nu(\mathbf{x}, \mathbf{u}) d\Omega \right] d\nu \quad (3)$$

This double integration is performed in the solver PRISSMA[1] using the discrete ordinate method for the angular integration associated with a spectral model [5, 7, 6].

2.2. The Discrete Ordinate Method (DOM)

The discrete ordinate method was originally proposed by Chandrasekhar [14] for astrophysical applications. The differential form of the RTE, Eq. (1),

[1] PRISSMA: Parallel RadIation Solver with Spectral integration on Multicomponent mediA, `http://www.cerfacs.fr/prissma`

is solved on a discrete set of directions of the solid angle as:

$$\int_{4\pi} f(\mathbf{u}) d\Omega \simeq \sum_{i=1}^{N_{\text{dir}}} w_i^a f(\mathbf{u}_i) \tag{4}$$

where $\mathbf{u}_i$ are the discrete ordinates and ω_i^a are the weights of the angular quadrature.

The spatial discretization is based on the Finite Volume Method (FVM) [15, 16, 17] and the Diamond Mean Flux Scheme (DMFS) is used for the spatial integration [8] on unstructured mesh. The angular quadrature is based on a set of N_{dir} directions (ordinates) using a S_n quadrature (with $N_{\text{dir}} = n(n+2)$) [18] or a LC_{11} quadrature (where $N_{\text{dir}} = 96$) [19].

2.3. Spectral Modelling

The spectral properties of absorbing gases such as CO, CO_2 and H_2O are known but not easy to handle. The spectroscopic data used here cover wavelengths in the range $\nu = [150; 9300]$ cm^{-1} and contain gas properties for 367 narrow bands of width $\Delta\nu_i = 25$ cm^{-1} [20]. Four additional bands $\nu = [9300; 20000]$ cm^{-1} are added to the visible spectrum to evaluate soot radiation using the correlation: $\kappa_{\nu,\text{soot}} = 5.5 f_v \nu$ (f_v is the volumic fraction of soot) [21].

Narrow-band models such as SNB-CK [22, 23] offer a good accuracy with a 5 points Gauss-Legendre quadrature and 371 bands. The spectral integration of the RTE in Eq. (1) leads to:

$$\int_0^\infty f_\nu d\nu \simeq \sum_{i=1}^{N_{\text{band}}} \sum_{j=1}^{N_{\text{quad}}} \omega_j \Delta\nu_i f_{i,j} \tag{5}$$

with $\Delta\nu_i$ being the width of the i^{th} band. This model leads to 1855 resolutions of the RTE per direction, which is too expensive to handle complex geometries in unsteady coupled calculations.

Global models are preferred (such as WSGG [24], SNB-FSK [25], SNB-FSCK [26]), which reduces the calculation to only 3 to 15 spectral integra-

tions, in Eq. (1), for each direction:

$$\int_0^\infty f_\nu d\nu \simeq \sum_{j=1}^{N_{\text{quad}}} \omega_j f_j \tag{6}$$

With the SNB-FSK and SNB-FSCK models, the absorption coefficient κ_j are calculated from narrow band properties. With these models the spectral calculation becomes more important than the spatial calculation so it is preferable to use tabulated coefficients. If the pressure is assumed constant, absorption coefficients can be tabulated in a four-dimensional space including temperature and H_2O, CO_2 and CO [7]. A sensitivity analysis of the ratio accuracy/CPU restitution time to the spectral model has been performed in previous work [7] and the retained spectral model consists on a tabulated SNB-FSK approach.

2.4. Sweeping order optimization

The spatial discretization is based on a finite volume method (FVM) with cell-centered scheme in an explicit solver. The balance of the radiation intensity at a given control volume depends on the flux crossing its faces and the intensity emitted by the volume. Each cell face is either an upstream or a downstream face for a given discrete direction. Sweeping the mesh in an arbitrary order may lead to undetermined upstream radiation intensities if the upstream cell was not previously resolved. This value is initially zero and the computation over all the mesh is restarted until the convergence of the solution.

An optimisation consists to define a sweeping order for each discrete direction. The sweeping order depends only on the chosen angular quadrature and the grid. The ordering is calculated by starting from a wall face where the upstream intensities are defined by the boundary conditions. The cell order is calculated such as the upstream intensity is known on the upstream faces from the previous cell in the discrete direction, see Fig. 2a.

The integration is performed only once in each cell by direction in the ordered mesh, *i.e.* no iterations are needed to converge the solution for a

problem without diffusion or reflection boundaries.

2.5. Reflecting boundaries

In the case of reflecting boundaries, the solution is calculated by an iterative procedure. The convergence criterion is evaluated as the relative error between two iterations on the sum of the wall incoming flux defined as:

$$H_\nu(\mathbf{r}_w) = \int_{2\pi} L_\nu(\mathbf{r}_w, \mathbf{u}) \cdot |\mathbf{n}_w \cdot \mathbf{u}| \, d\Omega \tag{7}$$

The total wall incoming flux on all boundaries for the convergence iteration n is defined as:

$$H_{\text{tot}}(n) = \int_{\text{Boundaries}} H_\nu(\mathbf{r}_w) d\mathbf{r}_w \tag{8}$$

the relative error is defined as:

$$\epsilon_H = \frac{|H_{\text{tot}}(n) - H_{\text{tot}}(n-1)|}{H_{\text{tot}}(n)} \tag{9}$$

The convergence criterion, Eq. 9, is evaluated at each discrete frequency $i_{\text{band}}, i_{\text{quad}}$ for the narrow band model, see Fig. 3b. For global models the convergence is evaluated on the spectral integrated quantity $H_{\text{tot}} = \int_0^\infty H_{\text{tot},\nu} d\nu$ which reduces the time calculation for the same accuracy, see Fig. 3c.

2.6. Numerical resolution

The radiation calculation has been coupled with LES calculations of unsteady turbulent combustion in previous works [5, 6]. The coupling methodology is presented in detail in [7]. From the radiative point of view, the general algorithm is summarized in Fig. 3a and link two main steps.

1. A preprocessing procedure generates the geometrical elements needed by the radiative calculation such as the mesh connectivities, the neighboring information and the sweeping ordering.

2. The radiative calculation is described. For an uncoupled calculation there is no coupling point and the initial physical solution is read in the preprocessing stage. For coupled simulations, the radiative calculation is

repeated at the different coupling points. Radiation is calculated every N_{it} fluid iterations, where N_{it} is the coupling frequency [7]. At each coupling point, the radiative solver receive an instantaneous solution with the physical fields from the fluid solver Φ: temperature, pressure, molar fraction of gases (H_2O, CO_2 and CO) and the soot volumic fraction. For each radiative calculation:

- In the case of a sub-domain decomposition, the mesh is partitioned with a standard recursive coordinate bisection method (RCB) (*i.e.* the mesh partitioning depends on the number of cores allocated for the radiative solver).

- The absorption coefficients are calculated with the chosen spectral model. For the narrow band model the frequency discretization rely on two indexes: the narrow band i_{band} and the quadrature point i_{quad} as in Eq. (5). For the global models the frequency discretization is given by the quadrature point over the full spectrum i_{quad} as in Eq. (6).

- The geometrical integration is then performed using the angular quadrature and the spatial scheme. At each discrete frequency the grid is scanned for each direction of the Eq. (4) using the sweeping order.

The measured CPU times in the rest of this paper concern the spectral and the geometrical calculations.

The physical fields $\Phi(T, P, X_{H_2O}, X_{CO_2}, X_{CO}, f_v)$ from the fluid solver are stored at the grid nodes Φ_j. In the case of global models the spectral calculation is performed at the grid nodes and then gathered at the cell center: $\kappa_j(\Phi_j)$, $L_j^0(T_j) \ \rightarrow \kappa_i, \ L_i^0$ (Fig. 2c). For the narrow band model the spectral calculation and the geometrical integration are calculated for each bands. It was preferred to gather physical fields $\Phi_i \rightarrow \Phi_j$ and calculating κ and L^0 at the cell to avoid multiple gathering operation on the bands.

The results are treated in a post-calculation stage at the end of each calculation. The radiative solution is scattered on the grid node (Fig.2d) and the solution is reconstructed on the whole domain for a sub-domains

decomposition. The radiative source term and the radiative heat flux are sent to the fluid solver at the end of each coupling point.

3. Parallelism strategies

3.1. Task decomposition: directions and frequencies

In the task decomposition, the calculation is split in independent parts over the cores and the total solution is obtained by addition of the contributions of each core. The solid angle and the gas spectra sumatories shown in Figs. 3b-3c and given in Eq.(4), (5) and (6), are used for the tasks decomposition.

However this kind of parallelism has strong limitations. The maximum scalability is fixed by the shape of the problem. Typically for combustion applications, radiative calculations rely on an S_4 angular quadrature (24 directions) and global spectral model (from 3 to 15 spectral points) so the maximum number of cores ranges from 72 to 360.

The memory use of the task decomposition depends on the numerical strategy and the CPU architecture. In a distributed memory architecture, MPI process runs on each core and the whole domain data is duplicated in the memory of each core. This approach can lead to a very high memory cost for complex configurations.

For shared memory architectures, each node of a cluster has N cores that share the same memory, it is more convenient to run the application using OpenMP. In this case the code runs N threads on the N cores of the node of the cluster using the common shared memory. The memory use by core decreases with the number of threads. The OpenMP task decomposition is performed over directions only, so $N_{\text{thread}} \leq N_{\text{dir}}$. If the number of directions is greater than the number of cores per node of the cluster, the MPI and the OpenMP decompositions can be combined. For example, in a calculation over 24 directions with 3 nodes of 8 cores, there should be 3 MPI process with 8 OpenMP threads per calculation node.

The the geometrical calculation presented in the Fig. 3a is detailed in Fig. 3b for the narrow band model and in Fig. 3c for the global models. At

each discrete frequency the grid is scanned for each direction of the Eq. (4) using the sweeping order given by the array *pathway*, see Figs. 3b-3c. At each convergence iteration, H (from Eq. 7) is integrated on the directions doing a reduction on the MPI process, see Figs. 3b-3c.

3.2. Sub-domain decomposition

The size of the calculation is reduced by splitting the physical space with the sub-domain decomposition. The memory usage and the CPU cost of each domain decrease with the size of the sub-domain.

To be efficient this kind of parallelism needs to solve independent parts like and there is not theoretical scalability limit. Practically the scalability is limited by a minimal size for the sub-domains due to the cost of the communication between the domains. For explicit CFD solver the sub-domain decomposition is very efficient with hundreds and thousands cores. For the radiative transfer the sub-domains are not independent due to non-local exchanges. This decomposition is *a priori* less efficient and communications between sub-domains must be incorporated.

A new sub-domain decomposition has been implemented in the solver PRISSMA using an iterative procedure to reduce the number of required communication. The faces between two sub-domains defined a"virtual boundary". This new boundary condition inside each domain is treated like a reflecting boundary using the iterative procedure to converge the solution presented in section 2.5.

In the first stage of the radiative calculation presented in Fig. 3a, the grid is split in N_{domains}, with $N_{\text{domains}} \leq$ total number of MPI process. The sweeping order array is also split for each direction over the sub-domains. If the adjacent given cell of the sweeping, belongs to another domain, the downstream face is marked as a "virtual face", see Fig.2b, defining virtual boundaries.

For the geometrical integration in Fig. 3a, the radiation intensity is stored at virtual boundary faces for each direction and frequency: $L_{\text{virt}}(i_{\text{dir}}, i_{\text{quad}}, i_{\text{band}})$. This array is global, defined on the full domain, and is updated at each convergence iteration with a reduction over all sub-domains, see Figs. 3b- 3c.

The total intensity on the virtual faces for the convergence iteration n is calculated as

$$L_{\text{virt,tot}}(n) = \int_{\text{Virtual Boundaries}} L_{\text{virt}}(\mathbf{x}) d\mathbf{x} \tag{10}$$

The convergence criterion is calculated on relative error between two iterations on the sum of L_{virt} as:

$$\epsilon_{L_{\text{virt}}} = \frac{|L_{\text{virt,tot}}(n) - L_{\text{virt,tot}}(n-1)|}{L_{\text{virt,tot}}(n)} \tag{11}$$

As for reflection the convergence is evaluated: at each discrete frequency for the narrow band model, after the spectral integration for the global models, see Figs. 3b and 3c. The final criterion is given by the maximal error of the reflection and the sub-domains tests:

$$\epsilon = \max(\epsilon_H; \epsilon_{L_{\text{virt}}}) \tag{12}$$

4. Results and discussions

4.1. Configuration and test case

The different parallel strategies have been tested with a cluster on which each node has:

- two processors AMD 12 cores 2.2 Ghz (4 flop per cycle per core *i.e.* 211 GFlops of peak performance by node),
- 32 Gb memory by node (DDR2 memory at 1333 MHz),
- 64 Kb cache L1 (instructions and data), 512 Kb cache L2 per core,
- 12 Mb cache L3 shared by 12 core of each processor.

The network Infiniband offers a band-width of 5 Gb/s between nodes with a MPI latency less than 1 μs. The installed operating system on the node is RedHat Entreprise 5.5.

The chosen application is a sector of an annular helicopter combustion chamber from Turbomeca, which was studied recently by Amaya *et al.* [27] and Staffelbach *et al.* [28, 2]. The operating point considered corresponds to full thrust. The radiative domain is extracted from the fluid domain with a mesh of 2.6M of cells. The initial solution is given on the Fig. 4.

The radiative calculation is performed with the DMFS scheme for the spatial integration along with a S_4 (24 directions) angular quadrature. The spectral integration uses the tabulated SNB-FSK model with 5 spectral quadrature points.

The boundary conditions are defined on the walls by an emissivity of 0.8 with temperature given by Turbomeca. On the periodic boundaries of the sector the boundary is set as purely reflecting wall.

The efficiency of the different parallelism decomposition is evaluated by regarding the CPU time reduction with the number of cores compared to the single core calculation. All tested decompositions are given in the Table 1. It was verified that the numerical solution for parallel calculation tends to the single processor calculation by adjusting the criterion error.

4.2. Task decompositions (A)

4.2.1. Directions (A1,A2)

The task decomposition on the directions is tested using the distributed memory (MPI process, A1) and the shared memory (OpenMP threads, A2) numeric strategies, see Fig. 5. The ideal CPU time reduction is calculated as the linear CPU time reduction with the number of cores. The two strategies give a good efficiency of the task decomposition on the directions.

The MPI case is closer to the ideal than the OpenMP case. With MPI, only the array H (Eq. 7) is reduced at the end of the direction loop for the reflection convergence criterion and all the MPI process contribute to the reduction with the instruction *MPI_ALL_REDUCE*. Others quantities such as the incident intensity or the radiative flux are reduced in the postprocessing stage. With OpenMP, all calculated quantities are reduced at the end of the directions loop by only one master thread while other threads are waiting so the efficiency slightly decreases.

4.2.2. Directions and frequencies (A3)

The direction/frequency decomposition (A3) presented in section 3.1 has been implemented using the MPI numeric strategy and was tested on a number or cores form 24 to 120 as shown on Fig. 5b.

The combination of the frequency and the direction decompositions based on MPI is completely inefficient after 24 cores when using the frequency decomposition and the calculation time stagnates. As the convergence is tested on the spectral integrated incident wall flux $H = \int H_\nu d\nu$, all the frequencies are not strictly independent and a MPI communications are needed at the end of the loop over directions.

4.3. Domain decomposition (B)

The domain decomposition (B) requires the addition of an iterative loop. At each sub-iteration the radiation intensity is updated on the virtual boundaries. Fig. 6 gives the number of convergence iteration with the number of sub-domains for a convergence criterion of 1%, Eq. 12. Single core computations need only three iterations to take into account the reflection on the periodic boundaries. There is an important jump with two domains reaching nine iterations to converge, then the number of iterations increases slowly: 0.114 additional iterations per domain up to 48 domains and 0.056 further on.

In a coupled simulation the radiative source term is computed at each coupling points (see Fig. 3a). The incident flux at the wall H and the radiation intensity on the virtual boundaries L_{virt} are stored between two coupling point in order to have a good initial prediction for the next iteration. A second coupling point contains the solution of the previous coupling point for the wall and the virtual boundaries. Fig. 6b shows the number of convergence sub-iterations for each coupling point a in coupled DOM/LES simulation with a domain decomposition on 8 domains. The first calculation needs 11 sub-iterations for a convergence criterion of 1% on the boundaries and virtual boundaries, the following coupling points converge in only 4 iterations thanks to the storage and reuse of the previous solution.

Two cases are distinguished to evaluate the efficiency of the sub-domains

decomposition: case B–i with a non-limited number of sub-iterations and case B–ii with a number of sub-iterations limited to 4 reproducing the coupling points after the first radiative calculation. The chosen number of 4 iterations ¡was obtained in the case of a coupled calculation with only 8 domains. A further investigation would be needed to evaluate the number of convergence iterations after the second coupling point, here it was assumed that this value was constant as the number of convergence iteration increases slowly with the number of domains, Fig. 6.

The reduction of the calculation time for the cases B–i and B–ii is plotted on the Fig. 7. Two ideal efficiencies are plotted as the linear reduction of the calculation time from the case with two domains for the cases B–i and B–ii.

The two sub-domains decomposition cases have similar trends: each curve fits well with the ideal cases. The domain decomposition is efficient from 2 to 24 sub-domains and the efficiency of the time reduction decreases, the CPU time is almost constant when increasing the number of domains. Figure 7b shows the number of grid nodes per sub-domain: after 48 cores the number of node per domains is less than 10^5. This value seems to be insufficient to get an efficient domain decomposition as the calculation time becomes too small compared to the communication time. The test configuration is relatively small and the domain decomposition should be tested on a larger configuration.

Figure 7 shows that the task decomposition A3 (direction/frequency) is more efficient than the domain decomposition for the case B–i. The overcost due to the additional convergence sub-iteration is not balanced with a better scalability. In the case B–ii, the domain decomposition is better than the task decomposition from 24 cores. The domain decomposition has a better efficiency than the task decomposition in the case of unsteady coupled simulations.

4.4. Hybrid decompositions (C)

Some hybrid calculations were tested combining the direction decomposition based on the OpenMP strategy and frequency as well as domain

decomposition based on the MPI strategy. In the following, the ideal case is calculated as the linear reduction of the time calculation from the single core case.

4.4.1. Directions/Frequencies (C1)

The MPI decomposition on the frequencies is combined to the OpenMP direction decomposition (C1). This decomposition is limited by the number of tasks on the directions and frequency with a maximal number of cores of 120 in the studied case. The decomposition C1 is compared to the A3 decomposition (directions/frequency based on MPI only) on the Fig. 5b. Compared to the directions/frequency based on MPI only the hybrid decomposition C1 gives an improved efficiency for the reduction of the computational time after 48 cores.

The normalized speed-up by frequency is calculated as:

$$\mathrm{S1} = \frac{T_{\mathrm{CPU}}(24 \text{ directions})}{T_{\mathrm{CPU}}(n)} \tag{13}$$

where n is the number of frequencies used for the task decomposition. The speed-up by frequency is calculated for the case A3 (direction–MPI/frequency) and for the case C1 (direction–OpenMP/frequency) and plotted on the Fig. 8.

The combination of OpenMP for the directions and MPI for the frequencies strongly improves the speed-up compared to the MPI-tasks decomposition only. With the MPI-tasks decomposition the speed-up decreases and this decomposition is inefficient.

4.4.2. Directions/Sub-domains (C2)

The sub-domains decomposition is combined to the OpenMP direction decomposition (C2), see Fig. 9. Like previously, two cases are regarded: a non-limited number of convergence iteration case C2–i and a case limited to 4 iterations C2–ii.

For these cases there is an overcost between 24 and 48 cores due to first the domain decomposition. After 48 cores this hybrid decomposition is very efficient, especially in the case C2–ii.

The hybrid decomposition C1 (direction–OpenMP / frequency–MPI) has the best efficiency but is limited to 120 cores as it is a task decomposition (*i.e.* 24 directions with 5 frequencies). The case C2–ii (direction–OpenMP / sub-domains with 4 sub-iterations) is as much efficient as the case C1 from 120 cores. For both cases C2–i and C2–ii (directions–OpenMP / sub-domains), the reduction of the time calculation is shifted from the ideal case due to the sub-iterations but the curve slopes follows the ideal case from 96 to 600 cores (*i.e.* 4 to 25 sub-domains). The efficiency after 600 cores (25 domains) because the domain decomposition is less efficient when the domains are too small. The CPU time calculation decreases up to 1200 cores but with a reduced efficiency for both cases C2–i and C2–ii.

There is a quasi-constant ratio between the calculation time and the ideal case while the domain decomposition is efficient up to $N_{\text{proc}} = 600$, *i.e.* $N_{\text{domain}} = 25$. The ratio $T_{\text{CPU}}/T_{\text{Ideal}}$ is plotted on the Fig. 10 for the both hybrid OpenMP / MPI-subdomains calculations C2–i and C2–ii. While the domain decomposition is efficient ($N_{\text{domains}} < 25$) this ratio is almost constant and seems to be linked to the number of convergence iterations for the both cases. For $N_{\text{domain}} > 25$ the ratio $T_{\text{CPU}}/T_{\text{Ideal}}$ increases rapidly as the domain decomposition efficiency decreases.

4.4.3. Directions/Frequencies/Sub-domains (C3)

The frequency decomposition based on MPI and the subdomains decomposition are combined to the OpenMP direction decomposition (C3).

For N_{core} from 1 to $N_{\text{dir}} = 24$, parallelism is based on the direction decomposition with OpenMP. For N_{core} from $N_{\text{dir}} = 24$ to $N_{\text{dir}} \times N_{\text{quad}} = 24 \times 5 = 120$, the direction decomposition is combined with the decomposition on the frequencies. For N_{core} above $N_{\text{dir}} \times N_{\text{quad}} = 120$ cores the domain decomposition is used. Theoretically this decomposition could increase the scalability of the hybrid case OpenMP / Sub-domains by a factor equal to the number of spectral quadrature point (*i.e.* $N_{\text{quad}} = 5$).

This three levels decomposition (C3) was tested only with a maximum of 1200 cores as shown on Fig. 11 for the cases C3–i (non-limited iterations) and C3–ii (four convergence iterations). There is an overcost due to the domain

decomposition between 120 and 240 cores. After 240 cores the slope of the calculation time reduction follows the ideal case. The hybrid cases direction–OpenMP / sub-domains C2–i and C2–ii recover the same behaviour of the only sub-domains cases B–i and B–ii as already shown on Figs. 7 and 8: the CPU time reduction after 24 cores gives a good efficiency with a maximum of 25 domains. In the decomposition C3 the reduction of the CPU time has the same slope as the ideal case after 1200 cores (*i.e.* 10 domains).

With this three levels decomposition a parallel efficiency is expected up to 3000 cores as the domain decomposition is efficient up to 25 domains. A reduction of the time calculation is expected up to 6000 cores (24 directions x 5 frequencies x 50 domains). A larger amount of cores would be required to evaluate the efficiency with this high number of calculation cores.

In the study case the coupled CPU time ratio between the fluid solver AVBP[2] and the radiative solver PRISSMA is $T_{\mathrm{AVBP}}/T_{\mathrm{PRISSMA}} = 15$ for the same number of cores and a coupling frequency of $N_{\mathrm{it}} = 100$ (*i.e.* one radiative calculation for 100 fluid iterations [7]). The fluid solver has shown a good performance up to 45 000 cores [29]. The scalability of the radiation solver with the decomposition C3 enables to run coupled simulations on massively parallel architectures. The best communication strategy to reach optimal interactions is under investigation, for example to avoid bottleneck communications between the 2 codes.

4.4.4. Speed-up of the sub-domains decomposition

The normalized speed-up by sub-domain is calculated as:

$$\mathrm{S2} = \frac{2 \times T_{\mathrm{CPU}}(2 \text{ domain})}{T_{\mathrm{CPU}}(n)} \tag{14}$$

where n is the number of domains used for the sub-domain decomposition. The speed-up by domain is calculated for different cases and plotted on Fig. 12 for: B–i/B–ii (domains only), C2–i/C2–ii (directions/domains) and C3–i/C3–ii (directions/frequencies/domains). For each decompositions, two

[2] `www.cerfacs.fr/cfd/avbp.html`

cases were tested: a non-limited of convergence iteration (i) and a limited case to $N_{\max} = 4$ sub-iterations (ii). The speed-up has been normalized by the time calculation with two domains corresponding to a different number of cores: two in cases B, 48 in cases C2 (2 x 24 directions) and 240 in cases C3 (2 x 24 directions x 5 frequencies). The speed-up S2 is regarded for the different domain decompositions:

- In the cases B, the sub-domain decomposition has a speed-up close to the ideal case from 2 to 25 domains. The speed-up increases until 72 domains for the cases B–i and B–ii with a reduced slope.

- In the cases C2, the speed-up is shifted under the ideal case but with the same slope until 24 domains, after the speed-up still increases but with a reduced slope. In the case C2–ii, the speed-up is upper the ideal case (*i.e.* the speed-up is normalized with the calculation on 24 cores with a non-perfect speed-up) demonstrating that in this case the domain decomposition is more efficient and improves the direction decomposition based on OpenMP. After 24 domains, the case C2–ii fits very well with the ideal case.

- In the cases C3, the number of tested domains is significantly reduced and a larger amount of cores would be needed to fully evaluate this decomposition. The case C3–i has a speed-up higher than the ideal case between 2 and 5 domains and then the slope of the curve follows the ideal case until 10 domains. The case C3–ii is slower than the ideal speed-up but with a constant slope suggesting that this decomposition could give an important parallel efficiency with a larger number of domains and cores.

4.5. Memory use

The memory use per core is plotted on Fig. 13 for each decomposition. The task decompositions based on MPI (A1 and A3) are the least efficient with a constant memory use per core as data is duplicated on each cores.

For the sub-domain decompositions (cases B) a master core which contains the information to reconstruct the solution over the full domain is

distinguished from the other cores called slaves. The memory of the master is more important because it contains the data needed to reconstruct the final solution on the whole domain. The memory usage decreases for 12 cores/domains for the master and 24 cores/domains for the slave. After 24 cores/domains the memory usage increases sharply.

The hybrid decomposition OpenMP / MPI-subdomain (cases C2) decreases the memory consumption up to 120 cores close to the ideal case calculated linearly from the single core calculation.

5. Conclusion

For simulations of unsteady combustion process that include radiation, different decompositions were investigated for parallel calculations using DOM for the radiative solver. First a tasks decomposition is tested using the discrete directions and frequencies with two different numeric strategies: MPI (distributed memory) and OpenMP (shared memory). A sub-domains decomposition was implemented in the radiative solver based on an iterative convergence procedure. Then hybrid decompositions were investigated. These partitioning methods were tested on a sector of an industrial combustion chamber.

Concerning the tasks decomposition it was shown that a decomposition only on the directions was more efficient using MPI than OpenMP.

When combining the directions and the frequencies decomposition the use of MPI only is completely inefficient for the frequencies. The use of OpenMP for the directions and MPI for the frequencies has shown an improved efficiency. However the tasks decomposition is limited to a maximal number of cores fixed by the size of the discretization: N_{dir} x N_{quad}, *i.e.* 120 in this case.

The sub-domains decompositions were investigated in two situations: a non coupled calculation where the number of convergence sub-iterations was not limited and a case reproducing a coupled calculation after the first coupling point. For the case of an uncoupled calculation, it was shown that there is an overcost using domain decomposition due to an higher number

of convergence iterations. This overcost is not balanced out with a better parallel efficiency. CPU time is in favor of the tasks decomposition method, up to 120 cores. The domain decomposition has a good parallel efficiency for a maximum of 25 domains. Above this point, the efficiency decreases because the sub-domains become too small and a bigger configuration would be needed. In the case of a coupled calculation, it was shown that the number of iterations after the first coupling point is strongly reduced with a storage and reuse the previous solution on the interfaces between the domains. The coupled case converges in only four iterations. In this case the sub-domain decomposition has a better efficiency than the tasks decomposition, based on MPI above 30 cores, even if the efficiency shows a decrease above 25 cores.

Different hybrid decompositions are proposed. The directions decompositions based on OpenMP is combined with sub-domains. A very good efficiency is obtained for the both uncoupled and coupled cases. The domains decomposition follows the slope of the ideal case until 600 cores and the CPU time decrease until 1200 cores. A three level decomposition is then regarded, combining the direction with OpenMP, the frequencies with MPI and the sub-domains. The slope of the reduction of the CPU time follows the ideal case up to 1200 cores. A good parallel efficiency is expected until 3000 cores, corresponding to 25 domains.

The different parallel decompositions were regarded in terms of memory usage. Tasks decomposition based on MPI is not interesting because all the data is duplicated in the memory for each cores. Domain decomposition shows a reduction of the memory use for 12 to 24 cores but increases above this value. The hybrid decompositions OpenMP / sub-domains are very interesting in terms of memory use and very close to the ideal case at 120 cores.

Compared to the previous works on parallel DOM radiative calculations it has been shown that the best strategy is to combine different parallelism levels. In this work an important gap was passed to reach massively parallel calculations for the radiative heat transfer. This work should be used in the future to run coupled calculations with LES on complex geometries, such as

360° annular chambers, which are extremely demanding in terms of HPC.

Acknowledgements

The work was supported by the Fondation de Recherche en Aéronautique et Espace (FRAE) in the project STRASS : http://strass.cerfacs.fr, project coordinator E. Riber. We acknowledge Isabelle d'Ast from CERFACS for her help on the use of the cluster. Thanks to F. André from CETHIL for spectrum illustration.

References

[1] G. Boudier, L. Gicquel, T. Poinsot, Combust. Flame 155 (2008) 196–214.

[2] G. Staffelbach, P. Wolf, R. Balakrishnan, A. Roux, T. Poinsot, in: CTR (Ed.), Proceedings of the Summer Program, pp. 259–269.

[3] N. Gourdain, L. Gicquel, M. Montagnac, O. Vermorel, M. Gazaix, G. Staffelbach, M. Garcia, J. Boussuge, T. Poinsot, Comput. Sci. Discovery 2 (2009) 015003.

[4] R. Gonçalves dos Santos, M. Lecanu, S. Ducruix, O. Gicquel, E. Iacona, D. Veynante, Combust. Flame 152 (2008) 387–400.

[5] J. Amaya, O. Cabrit, D. Poitou, B. Cuenot, M. E. Hafi, J. Quant. Spectrosc. Radiat. Transf. 111 (2010) 295–301.

[6] D. Poitou, J. Amaya, M. El Hafi, B. Cuénot, Combust. Flame 159 (2012) 1605–1618.

[7] D. Poitou, M. El Hafi, B. Cuenot, J. Heat Transfer 133 (2010) 062701–10.

[8] D. Joseph, M. E. Hafi, R. Fournier, B. Cuenot, Int. J. Therm Sci. 44 (2005) 851–864.

[9] G. Krishnamoorthy, R. Rawat, P. J. Smith, Numer. Heat. Tran., Part B 47 (2004) 19–38.

[10] J. Gonçalves, P. J. Coelho, Numer. Heat. Tran., Part B 32 (1997) 151–173.

[11] S. D. Pautz, Nucl. Sci. Eng. 140 (2002) 111–136.

[12] S. J. Plimpton, B. Hendrickson, S. P. Burns, W. M. Iii, Nucl. Sci. Eng. 150 (2005) 267–283.

[13] G. Colomer, R. Borrell, O. Lehmkuhl, A. Oliva, in: 14 International Heat Transfer Conference.

[14] S. Chandrasekhar, Radiative Transfer, Dover, Mineola, NY, 1960.

[15] D. Joseph, P. J. Coelho, B. Cuenot, M. E. Hafi, in: Proceedings of Eurotherm73 – Computational Thermal Radiation in Participating Media.

[16] K. A. Jensen, J. Ripoll, A. Wray, D. Joseph, M. E. Hafi, Combust. Flame 148 (2007) 263–279.

[17] D. Joseph, P. Perez, M. E. Hafi, B. Cuenot, J. Heat Transfer 131 (2009) 052701–9.

[18] Truelove, Discrete-ordinate solutions of the radiation transport equation, J. S. (Univ. of Newcastle), 1987.

[19] R. Koch, R. Becker, J. Quant. Spectrosc. Radiat. Transf. 84 (2004) 423–435.

[20] A. Soufiani, J. Taine, Tech. Notein Int. J. Heat Mass Tran. 40 (1997) 987–991.

[21] F. Liu, H. Guo, G. J. Smallwood, M. El Hafi, J. Quant. Spectrosc. Radiat. Transf. 84 (2004) 501–511.

[22] V. Goutiere, F. Liu, A. Charette, J. Quant. Spectrosc. Radiat. Transf. 64 (2000) 299–326.

[23] V. Goutière, A. Charette, L. Kiss, Numer. Heat. Tran., Part B 41 (2002) 361–381.

[24] A. Soufiani, E. Djavdan, Combust. Flame 97 (1994) 240 – 250.

[25] D. Poitou, J. Amaya, C. Bushan Singh, D. Joseph, M. E. Hafi, B. Cuenot, in: Proceedings of Eurotherm83 – Computational Thermal Radiation in Participating Media III.

[26] F. Liu, M. Yang, G. Smallwood, H. Zhang, in: Proceedings of ICHMT, RAD04, Istanbul, Turkey.

[27] J. Amaya, E. Collado, B. Cuenot, T. Poinsot, in: Proceedings of the Summer Program, Center for Turbulence Research, NASA AMES, Stanford University, USA.

[28] G. Staffelbach, L. Gicquel, G. Boudier, T. Poinsot, Proc. Combust. Inst. 32 (2009) 2909–2916.

[29] L. Y. Gicquel, N. Gourdain, J. Boussuge, H. Deniau, G. Staffelbach, P. Wolf, T. Poinsot, Compt. Rendus Mec. 339 (2011) 104–124.

List of Tables

1 Different parallel decompositon from 1 to 1200 cores combining directions, frequencies and sub-domains. The cases B, C2 and C3 are tested for two situations: a non-limited number of sub-iterations (i) and a limited number of four sub-iterations (ii). √: tested, –: not tested, ×: not possible. 28

List of Figures

1 Tasks decomposition: **(a)** on the direction, **(b)** on the frequency. **(c)** Data decomposition on sub-domains 29

2 **(a)** Sweeping ordering optimization on the whole domain for the discrete direction $\mathbf{s}_i$ **(b)** Splitting of the sweeping order optimisation on three sub-domains for the discrete direction $\mathbf{s}_i$, the dashed lines correspond to the virtual faces. **(c)** and **(d)** respectively Gather/Scatter procedure in the radiative solver. 30
3 Algorithm for the radiative calculation for **(a)** the general case of coupled calculation with a fluid solver; for the geometrical integration with the narrow band model **(b)** and the global models **(c)**. 31
4 Test case configuration, instantaneous solution from the fluid solver. Normalized fields for the temperature and the molar fraction of H_2O, CO_2 and CO. 32
5 Reduction of the time calculation with the number of processors **(a)** for the direction decomposition with the MPI (A1) and the OpenMP implementations (A2); **(b)** for direction/frequency tasks decompositions: (A3) MPI only and (C1) Hybrid OpenMp (directions) / Frequencies. 33
6 Number of convergence iterations **(a)** with the number of subdomains for a convergence criterion of 1%; **(b)** at different coupling points with 8 domains. 34
7 **(a)** Reduction of the time calculation with the number of processors for the task decomposition (A3, from 1 to 24 on directions, from 48 to 120 on frequencies) and the domains decomposition (B). Two cases are distinguished for the domain decomposition: a non limited case on the number of convergence iterations (B-i) and a limited case to $N_{\max} = 4$ iterations (B-ii). **(b)** Number of nodes per domain with the number of domains. 35
8 Normalized speed-up of the frequencies decomposition for the tasks decomposition (A3) and for the hybrid Directions (OpenMP) / Frequencies (C1). The speed-up is normalized by the calculation decomposed on 24 directions. 36

9 Reduction of the time calculation with the number of processors for hybrid decomposition: OpenMp (directions) / Frequencies (C1) and Directions (OpenMP) / Domains (C2). The hybrid calculation with the subdomains is run for 2 cases: a non limited case on the number of convergence iterations (C2-i) and a limited case to $N_{\max} = 4$ iterations (C2-ii). . . . 37
10 Ratio of the computational time and the ideal time for the hybrid calculation Directions (OpenMP) / Domains for 2 cases: a non limited case on the number of convergence iteration (C2-i) and a limited case to $N_{\max} = 4$ iterations (C2-ii). This ratio is compared to the number of convergence iterations for the two cases. 38
11 Reduction of the time calculation for the three levels decomposition Directions (OpenMP) / Frequencies / Domains (C3) for 2 cases: a non limited case on the number of convergence iteration (C3-i) and a limited case to $N_{\max} = 4$ iterations (C3-ii). 39
12 Normalized speed-up of the sub-domains decomposition in the different cases: sub-domains only (B), Directions (OpenMP) / Sub-domains (C2) and Directions (OpenMP) / Frequencies / Sub-domains (C3). For each decomposition a non limited case on the number of convergence iteration (i) and a limited case to $N_{\max} = 4$ iterations (ii). 40
13 Relative memory usage by processor for the MPI-Tasks, the sub-domains and the hybrid Direction (OpenMP) / Domains decomposition. The reference memory usage is 1.47Gb. . . . 41

	Number of cores:	1→ 24	→ 120	→ 1200
A1	Dir. (MPI)	√	×	×
A2	Dir. (OpenMP)	√	×	×
A3	Dir. (MPI) / Freq.	√	√	×
B	Domains	√	√	–
C1	Dir. (OpenMP) / Freq.	√	√	×
C2	Dir. (OpenMP) / Domains	√	√	√
C3	Dir. (OpenMP) / Freq. / Domains	√	√	√

Table 1: Different parallel decompositon from 1 to 1200 cores combining directions, frequencies and sub-domains. The cases B, C2 and C3 are tested for two situations: a non-limited number of sub-iterations (i) and a limited number of four sub-iterations (ii). √: tested, –: not tested, ×: not possible.

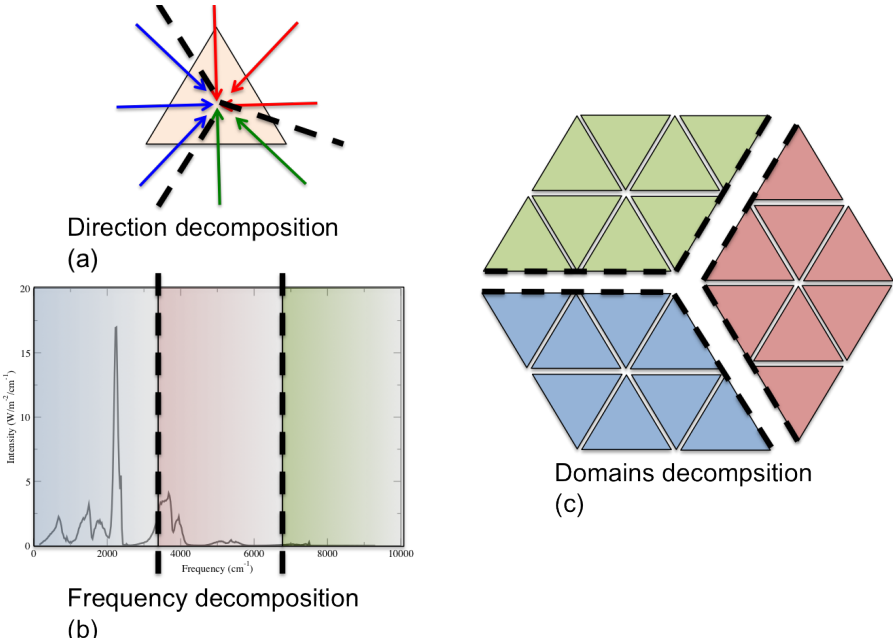

Figure 1: Tasks decomposition: **(a)** on the direction, **(b)** on the frequency. **(c)** Data decomposition on sub-domains .

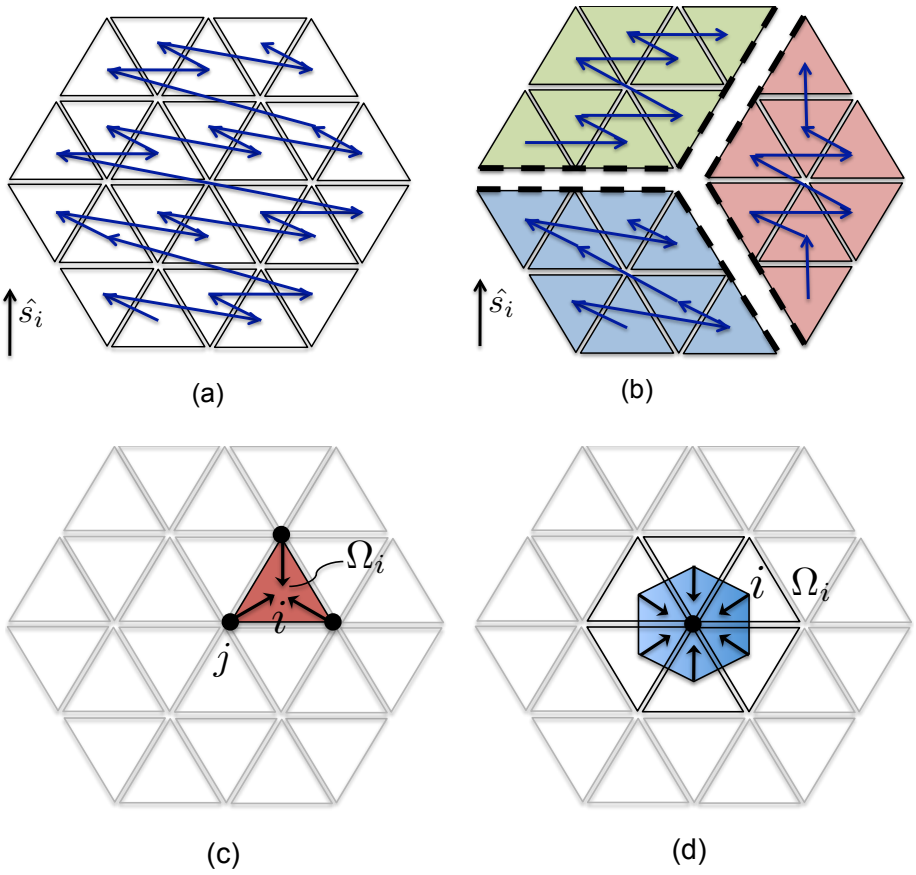

Figure 2: **(a)** Sweeping ordering optimization on the whole domain for the discrete direction $\mathbf{s}_i$ **(b)** Splitting of the sweeping order optimisation on three sub-domains for the discrete direction $\mathbf{s}_i$, the dashed lines correspond to the virtual faces. **(c)** and **(d)** respectively Gather/Scatter procedure in the radiative solver.

PRE PROCESSING

• Building geometrie and sweeping

RADIATIVE CALCULATION

• Mesh partitionning

Do i_{coupl} = 1, N_{coupl}

• Get fluid flieds from the fluid solver

$$\Phi = (T, P, X_{CO_2}, X_{H_2O}, X_{CO}, f_{v,soot})$$

Spectral Calculation

Angular integration

• Post-calculation of the results

• Send S_r and Q_w to the fluid solver

•EndDo

(a)

Do i_{band} = 1, N_{bands}
  Do i_{quad} = 1, N_{quad}
    WHILE CONVERGENCE > 1% Do
      Do i_{dir} = 1, N_{dir}
        Do i_{cell} = 1, N_{cell}(domain)
          X = pathway (i_{cell}, i_{dir})
          $S_r(\mathbf{x}) \simeq \sum_{j=1}^{N_q}\sum_{j=1}^{N_q} \omega_j \kappa_{i,j} \left(4\pi\sigma L^0_{i,j}(\mathbf{x}) - \sum_{k=1}^{N_{dir}} \omega^a_k L_{i,j}(\mathbf{x}, \mathbf{u}_k)\right)$
          Update H and L_{virt}(i_{dir},i_{quad},i_{band})
        EndDo
      EndDo
      MPI_REDUCTIONS: H (local subdomain), L_{virt}(all domains)
      TEST CONVERGENCE = max(REFLECTION, SUBDOMAINS)
    END WHILE
  EndDo
EndDo

(b)

WHILE CONVERGENCE > 1% Do
  Do i_{quad} = 1, N_{quad}
    Do i_{dir} = 1, N_{dir}
      Do i_{cell} = 1, N_{cell}(domain)
        X = pathway (i_{cell}, i_{dir})
        $S_r(\mathbf{x}) \simeq \sum_{j=1}^{N_q} \omega_j \kappa_i \left(4\pi\sigma T^4(\mathbf{x}) - \sum_{k=1}^{N_{dir}} \omega^a_k L_j(\mathbf{x}, \mathbf{u}_k)\right)$
        UPDATE H and L_{virt}(i_{dir},i_{quad},i_{band})
      EndDo
    EndDo
  EndDo
  MPI_REDUCTIONS: H (local subdomain), L_{virt}(all domains)
  TEST CONVERGENCE = max(REFLECTION, SUBDOMAINS)
END WHILE

(c)

Figure 3: Algorithm for the radiative calculation for **(a)** the general case of coupled calculation with a fluid solver; for the geometrical integration with the narrow band model **(b)** and the global models **(c)**.

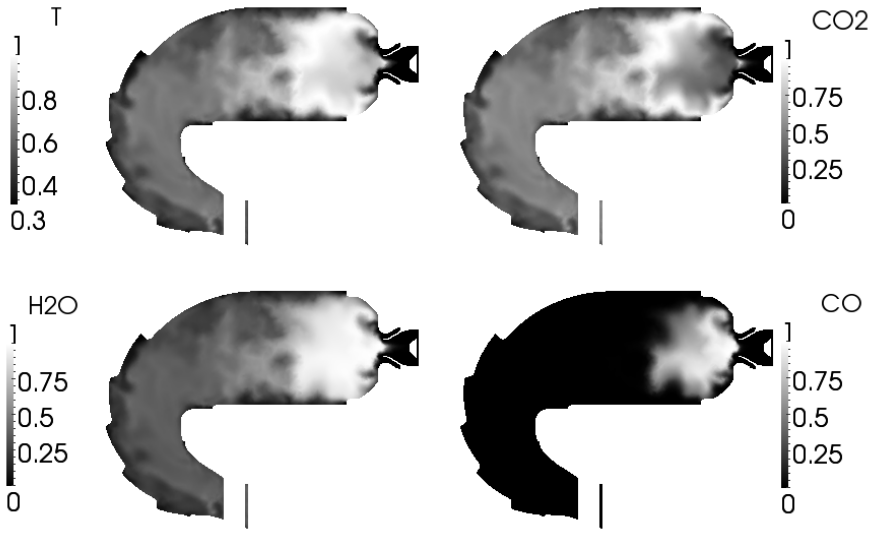

Figure 4: Test case configuration, instantaneous solution from the fluid solver. Normalized fields for the temperature and the molar fraction of H_2O, CO_2 and CO.

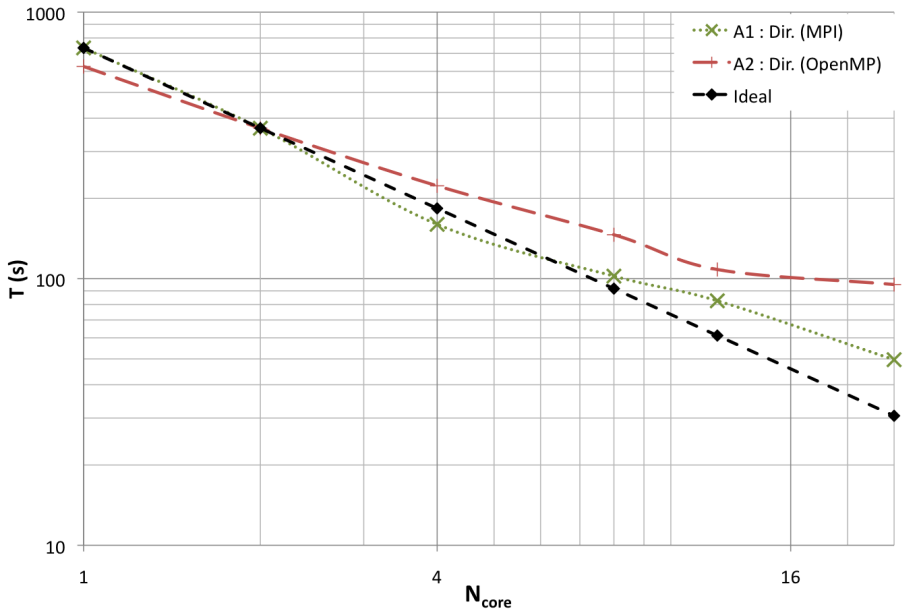

(a)

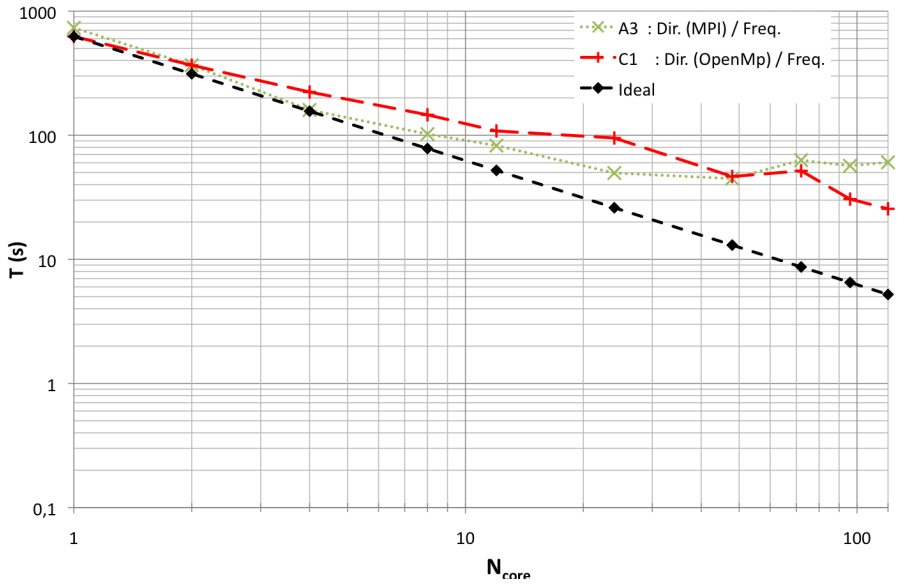

(b)

Figure 5: Reduction of the time calculation with the number of processors **(a)** for the direction decomposition with the MPI (A1) and the OpenMP implementations (A2); **(b)** for direction/frequency tasks decompositions: (A3) MPI only and (C1) Hybrid OpenMp (directions) / Frequencies.

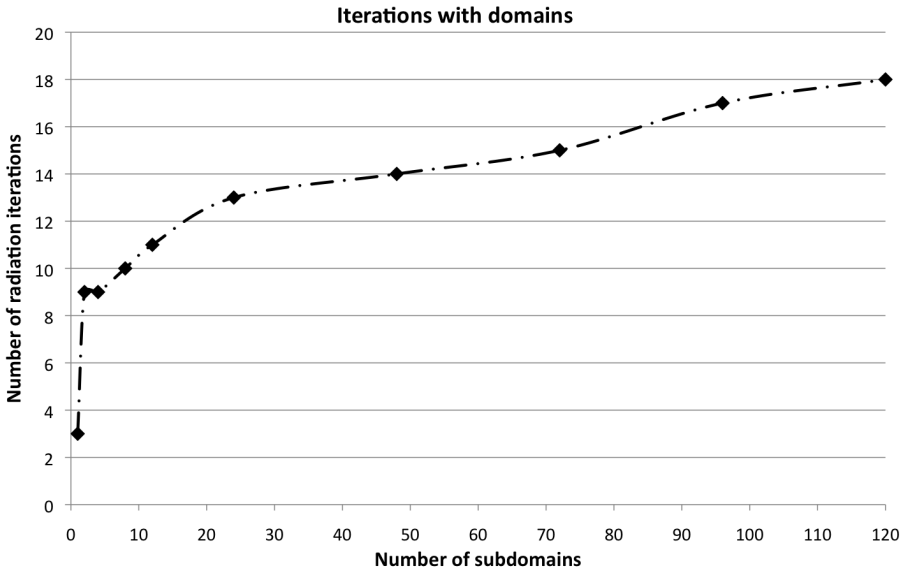

(a)

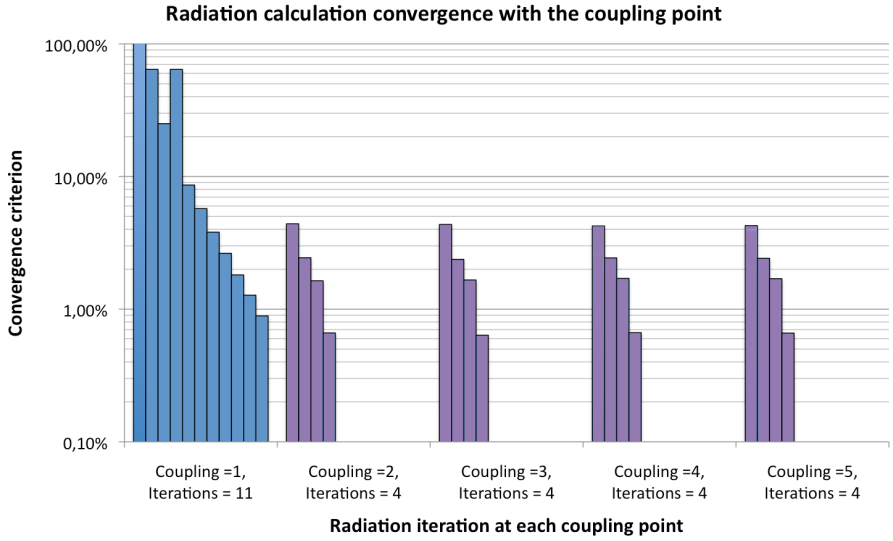

(b)

Figure 6: Number of convergence iterations **(a)** with the number of sub-domains for a convergence criterion of 1%; **(b)** at different coupling points with 8 domains.

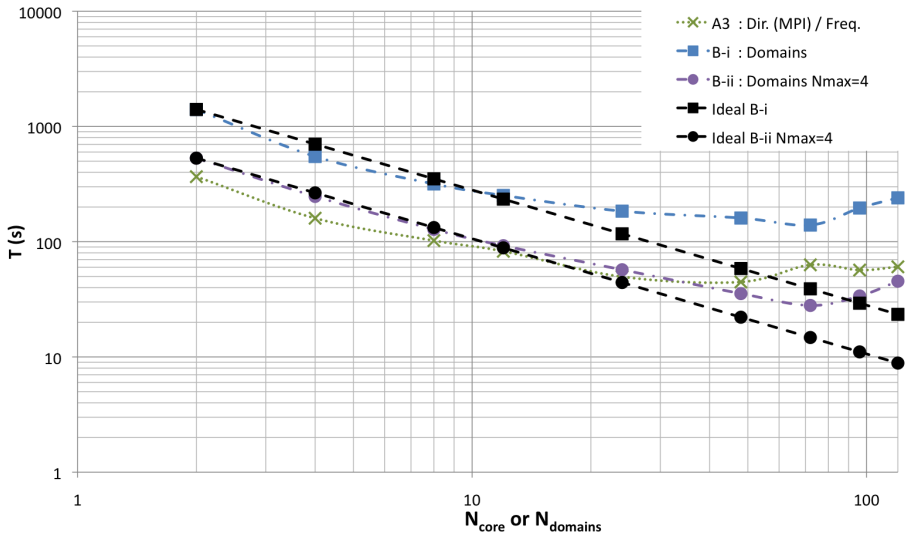

(a)

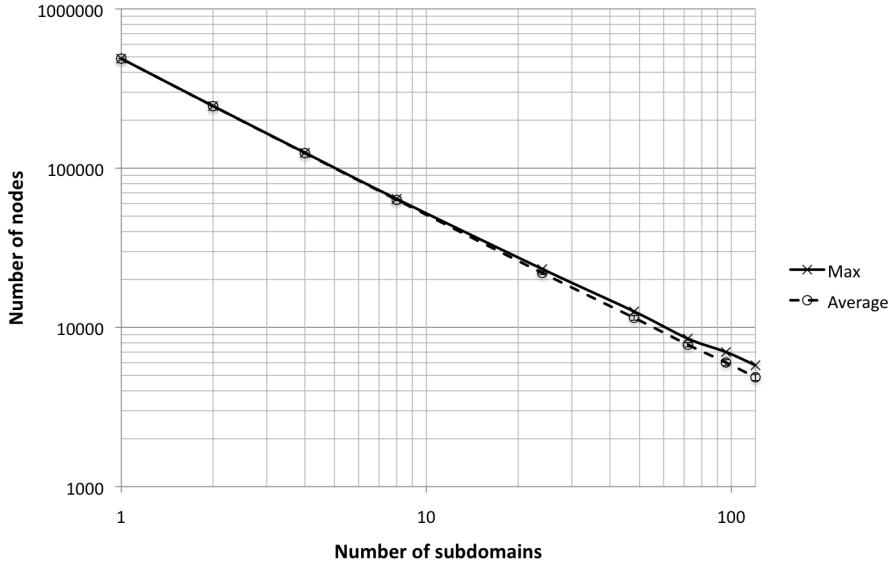

(b)

Figure 7: **(a)** Reduction of the time calculation with the number of processors for the task decomposition (A3, from 1 to 24 on directions, from 48 to 120 on frequencies) and the domains decomposition (B). Two cases are distinguished for the domain decomposition: a non limited case on the number of convergence iterations (B-i) and a limited case to $N_{\max} = 4$ iterations (B-ii). **(b)** Number of nodes per domain with the number of domains.

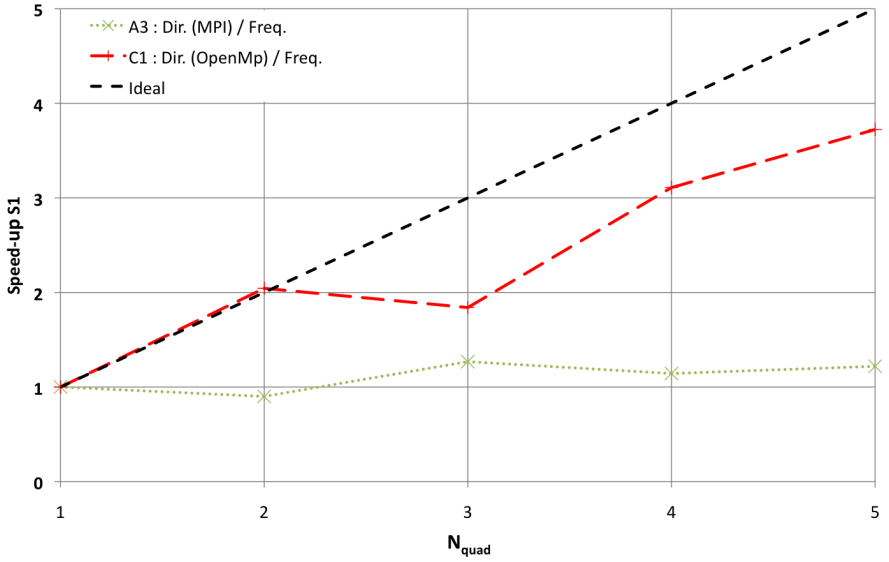

Figure 8: Normalized speed-up of the frequencies decomposition for the tasks decomposition (A3) and for the hybrid Directions (OpenMP) / Frequencies (C1). The speed-up is normalized by the calculation decomposed on 24 directions.

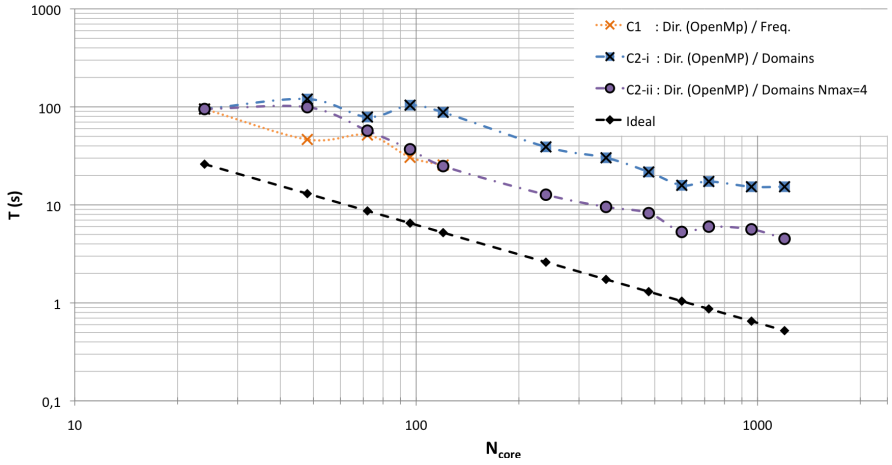

Figure 9: Reduction of the time calculation with the number of processors for hybrid decomposition: OpenMp (directions) / Frequencies (C1) and Directions (OpenMP) / Domains (C2). The hybrid calculation with the subdomains is run for 2 cases: a non limited case on the number of convergence iterations (C2-i) and a limited case to $N_{\max} = 4$ iterations (C2-ii).

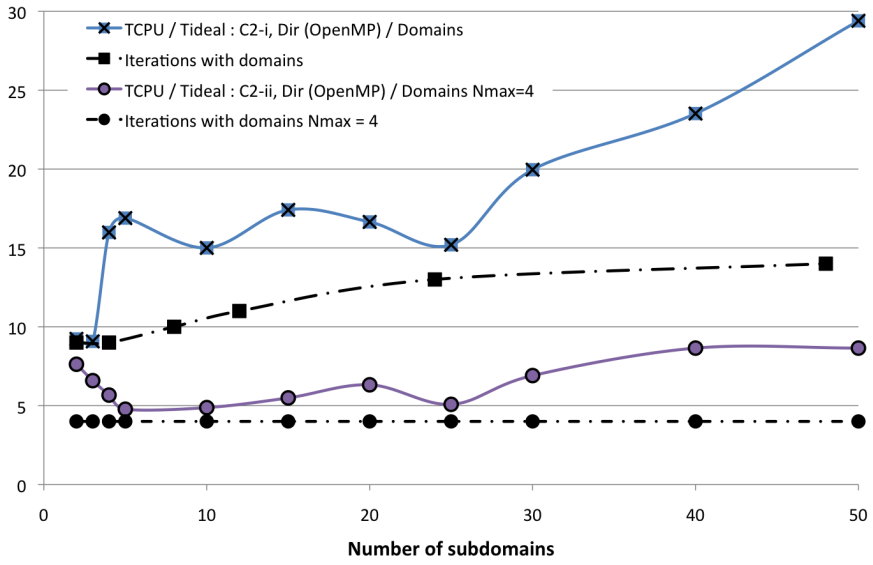

Figure 10: Ratio of the computational time and the ideal time for the hybrid calculation Directions (OpenMP) / Domains for 2 cases: a non limited case on the number of convergence iteration (C2-i) and a limited case to $N_{\max} = 4$ iterations (C2-ii). This ratio is compared to the number of convergence iterations for the two cases.

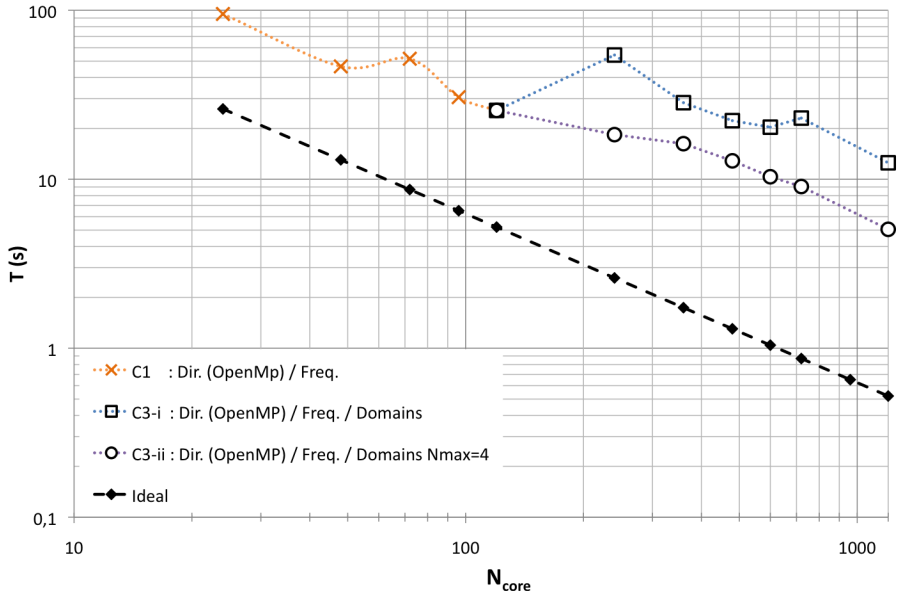

Figure 11: Reduction of the time calculation for the three levels decomposition Directions (OpenMP) / Frequencies / Domains (C3) for 2 cases: a non limited case on the number of convergence iteration (C3-i) and a limited case to $N_{\max} = 4$ iterations (C3-ii).

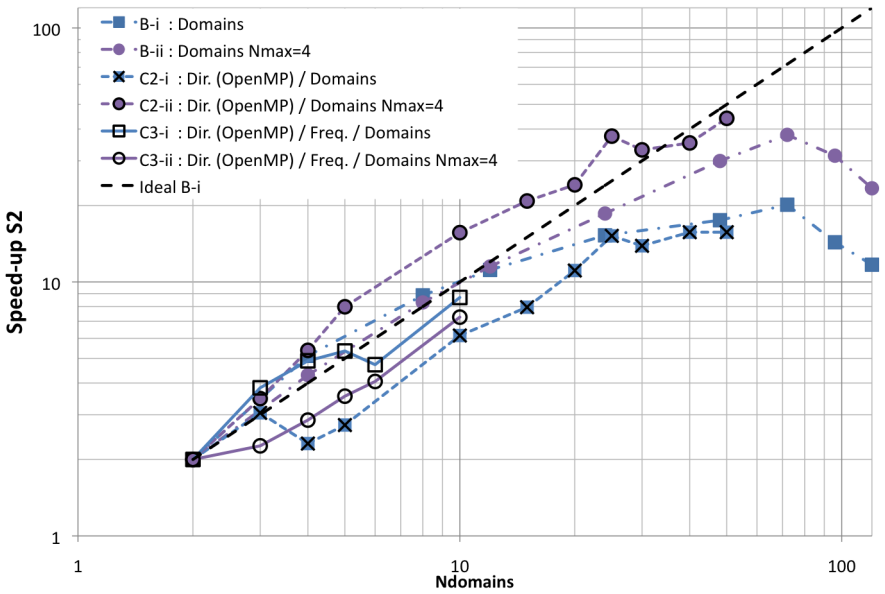

Figure 12: Normalized speed-up of the sub-domains decomposition in the different cases: sub-domains only (B), Directions (OpenMP) / Sub-domains (C2) and Directions (OpenMP) / Frequencies / Sub-domains (C3). For each decomposition a non limited case on the number of convergence iteration (i) and a limited case to $N_{\max} = 4$ iterations (ii).

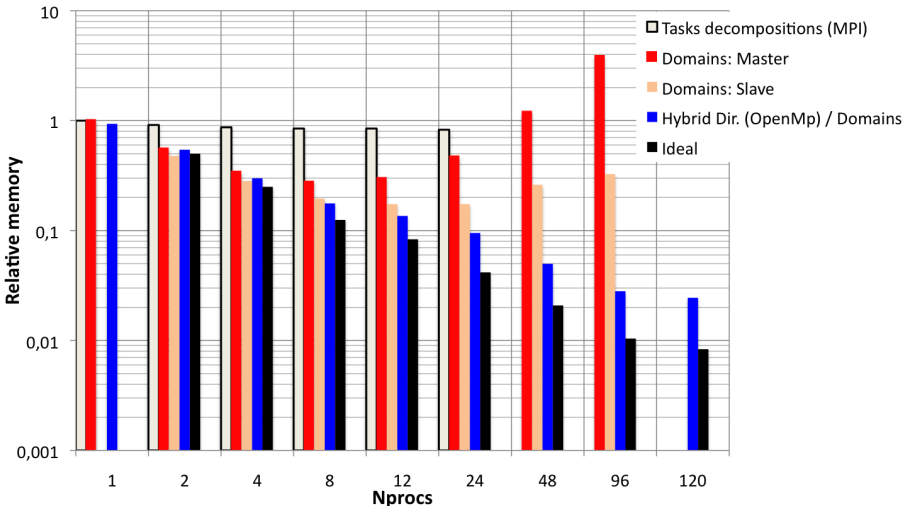

Figure 13: Relative memory usage by processor for the MPI-Tasks, the sub-domains and the hybrid Direction (OpenMP) / Domains decomposition. The reference memory usage is 1.47Gb.