

Alya: computational solid mechanics for supercomputers

E. Casoni · A. Jérusalem ·
C. Samaniego · B. Eguzkitza · P.
Lafortune · D. D. Tjahjanto · X. Sáez ·
G. Houzeaux · M. Vázquez

Received: date / Accepted: date

Abstract While solid mechanics codes are now conventional tools both in industry and research, the increasingly more exigent requirements of both sectors are fuelling the need for more computational power and more advanced algorithms. For obvious reasons, commercial codes are lagging behind academic codes often dedicated either to the implementation of one new technique, or the upscaling of current conventional codes to tackle massively large scale computational problems. Only in a few cases, both approaches have been followed simultaneously. In this article, we present a solid mechanics simulation strategy for parallel supercomputers based on a hybrid approach. Hybrid parallelization exploits the thread-level parallelism of multicore architectures, combining MPI tasks with OpenMP threads. This paper describes the proposed strategy, programmed in Alya, a parallel multi-physics code. Hybrid parallelization is specially well suited for the current trend of supercomputers, namely large clusters of multicores. The strategy is assessed through transient non-linear solid mechanics problems, both for explicit and implicit schemes, running on

E. Casoni, C. Samaniego, B. Eguzkitza, P. Lafortune, X. Sáez, G. Houzeaux and M. Vázquez
Barcelona Supercomputing Center (BSC-CNS), Edificio NEXUS I, Campus Nord UPC,
Gran Capitán 2-4, 08034, Barcelona, Spain
E-mail: mariano.vazquez@bsc.es

M. Vázquez
Artificial Intelligence Research Institute CSIC (IIIA-CSIC), Campus de la UAB, 08193 Bellaterra, Spain

A. Jérusalem and D. D. Tjahjanto
previous affiliation: IMDEA Materials Institute, Tecnogetafe, C/ Eric Kandel 2, 28906 Getafe, Spain

D. D. Tjahjanto
Department of Solid Mechanics, KTH Royal Institute of Technology, Teknikringen 8D, 10044 Stockholm, Sweden

A. Jérusalem
Department of Engineering Science, University of Oxford, Parks Road, Oxford, OX1 3PJ, UK
E-mail: antoine.jerusalem@eng.ox.ac.uk

thousands of cores. In order to demonstrate the flexibility of the proposed strategy under advance algorithmic evolution of computational mechanics, a non-local parallel overset meshes method (Chimera-like) is implemented and the conservation of the scalability is demonstrated.

Keywords Computational Mechanics · Finite Element Method · Parallel computing · Chimera

1 Introduction

With the increasing need for more advanced modeling techniques involving non-local approaches or multiphysics interactions, Finite Element (FE) solid mechanics codes requirements have traditionally slowed down the parallel efforts aimed at increasing the computational scalability of such codes. One of the direct consequences of this is that most commercial codes cannot efficiently scale in parallel computers when more than hundreds of cores are used. Academic FE codes, on the other hand, have often relied on the need to develop one unique technique of interest, potentially followed by a secondary development phase aimed at scaling it up because of the prohibitive cost of the technique. As a direct consequence, researchers have often focused on the scalability of one technique independently of the other ones; see for example the Arbitrary Lagrangian-Eulerian (ALE) methods [21], Discontinuous Galerkin (DG) methods [29,10], Fluid-Structure Interaction (FSI) [12,27,21], GPU solver [22], or even expensive constitutive models (and their also expensive meshing requirements) such as crystal plasticity [28]. It is also noticeable that the range of fields of application for these references is extremely wide, ranging from bio-medical, military, seismic, or fracture mechanics to polycrystalline texture analysis.

Large scale solid mechanics FE computation became a field of research by itself as soon as in the early 80's [16]. However, if preconditionner optimization studies for large scale computing were tackled shortly after [19], and barring a few exceptions (*e.g.*, Ref. [15] focusing on XFEM), they have almost always been focused on conventional FE codes (and often for one unique application) without consideration of all the previously listed evolutions [30,25,9,14]. As a result, many advanced codes are often optimized for large scale simulations of one application or technique, instead of being developed from the beginning for *any* large scale application, and then enhanced with advanced techniques following the same large scale framework imposed by the code structure.

In a somewhat opposite trend, fluid mechanics computational codes have generally been exploiting advanced techniques conjointly with parallel efficiency (see for instance OpenFOAM [5], CODE_SATURN [3]). This can be explained by the fact that fluid problems have traditionally required simpler and larger meshes than solid mechanics problems but also by the fact that constitutive model requirements might be much lighter for fluids than for solids. Among these codes, Alya [2] has been conceived as a computational mechanics code developed at Barcelona Supercomputing Center (BSC-CNS) aimed at

solving Partial Differential Equations (PDEs) in non-structured meshes. Alya exploits the similarities of the PDE-governed problems to solve with high parallel performance compressible and incompressible flows, thermal flows, excitable media or quantum mechanics for transient molecular dynamics [37, 18, 17, 13], among others, running in thousands of cores. Parallelization is hidden behind a common solver that assembles matrices and residuals and carries out the solution scheme. The scalability of the code thus exclusively depends on one unique set of parallel communication subroutines independently of the physics of the problem. Additionally, Alya is specially well-suited for the parallel solution of coupled multi-physics problems.

In the present work, we introduce the solid mechanics module of Alya. Using the large scale PDE solver capability of Alya's kernel, the solid module was implemented so that any future development of more complex FE techniques should conserve the high scalability of the code. These developments have been carried out along two parallel strategies: first, the Message Passing Interface (MPI) implementation, where the parallelization is based on a substructuring technique and uses a Master-Worker strategy [18]; and secondly, the OpenMP implementation used to treat many-core processors. The overall code is thus hybrid MPI/OpenMP. Due to their scalability, iterative schemes for sparse algebra with diagonal preconditioners are preferred. Non-symmetric problems are solved using GMRes or Bi-CGSTab schemes and symmetric ones, solved using Conjugate Gradient (CG) methods [31].

The Alya framework is specially well-suited to solve coupled problems. A Chimera overset meshes strategy is a particular way of coupling problems: two or more problems are solved in different meshes which are connected using their respective geometrical information. This connection scheme results in the fusion of all the problems in one single matrix. The problem can then be solved using the same parallelization scheme proposed before [13]. This scheme was thus implemented in Alya as an illustration of its large scale parallelization strategy where the code structure naturally lends itself to complex non-local approaches without loss of scalability performance.

The article is structured as follows. Section 2 introduces the governing equations and briefly describes the numerical method. Section 3 describes the structure of Alya, with special attention on the parallel strategy and the solver in Section 4. The parallelization and hybrid strategies are formally presented in Section 5 and the optimal scalability of the code is also shown. Section 6 presents some problems of applications where the Alya modules have been used with good performance. The paper ends with some conclusions and future lines.

2 Numerical scheme

The computational solid mechanics problem is solved using a standard Galerkin method for a large deformation framework using a generalized Newmark time

integration scheme. This framework, developed in a total Lagrangian formulation, is only briefly summarized below; for more details, see Ref. [24].

2.1 Standard Galerkin governing equations

Let $\boldsymbol{\varphi} : \mathbb{R}^3 \rightarrow \mathbb{R}^3$ be the function that maps a material point of an undeformed body $\mathbf{X} \in B_0$ in the reference configuration to its point $\mathbf{x} = \boldsymbol{\varphi}(\mathbf{X}) \in B$ in the current (or deformed) configuration. The deformation gradient tensor \mathbf{F} is defined as $\mathbf{F} := \nabla_0 \mathbf{x}$, where ∇_0 is the gradient operator with respect to the reference configuration. In cartesian basis, the components of \mathbf{F} are given by $F_{i,J} = \partial x_i / \partial X_J$. Index i in vector \mathbf{x} and J in vector \mathbf{X} are written in lower and upper case to indicate that the component is referred to in the current and reference configurations, respectively. Since $\mathbf{x}(\mathbf{X}) = \mathbf{X} + \mathbf{u}(\mathbf{X})$, where \mathbf{u} is the displacement vector, the deformation gradient can be given by $\mathbf{F} = \mathbf{I} + \nabla_0 \mathbf{u}$, with $\nabla_0 \mathbf{u}$ as the displacement gradient.

The equation of balance of momentum with respect to the reference configuration can be written as

$$\text{Div } \mathbf{P} + \mathbf{b}_0 = \rho_0 \ddot{\mathbf{u}}, \quad \forall \mathbf{X} \in B_0, \quad (1)$$

where ρ_0 is the mass density (with respect to the reference volume) and Div is the divergence operator with respect to the reference configuration, with $\text{Div } \mathbf{P} = \nabla_0 \cdot \mathbf{P}$. Tensor \mathbf{P} and vector \mathbf{b}_0 stand for the first Piola–Kirchhoff stress and the distributed body force on the undeformed body, respectively. Prescribed displacements and tractions are applied at reference boundary $\Gamma_0 = \Gamma_{d0} \cup \Gamma_{n0}$, where Γ_{d0} and Γ_{n0} correspond to the Dirichlet and Neumann boundary conditions, respectively:

$$\mathbf{u} = \bar{\mathbf{u}}, \quad \forall \mathbf{X} \in \Gamma_{d0}, \quad (2)$$

$$\mathbf{P} \cdot \mathbf{N}_0 = \bar{\mathbf{t}}_0, \quad \forall \mathbf{X} \in \Gamma_{n0}. \quad (3)$$

where \mathbf{N}_0 is the normal to the boundary in the reference configuration.

2.2 Space discretization

The weak form of balance of the momentum (1) can be formulated for any arbitrary admissible virtual displacement \mathbf{w} , such that,

$$\int_{B_0} \text{Div } \mathbf{P} \cdot \mathbf{w} \, dV + \int_{B_0} \mathbf{b}_0 \cdot \mathbf{w} \, dV = \int_{B_0} \rho_0 \ddot{\mathbf{u}} \cdot \mathbf{w} \, dV. \quad (4)$$

For a FE approximation $\Omega_0 = \bigcup_e \Omega_0^e$ of the undeformed continuum body B_0 , let \mathbf{u}_h be a polynomial approximation of degree k of the actual displacement \mathbf{u}

$$\mathbf{u}_h(\mathbf{X}) = \sum_a N^a(\mathbf{X}) \mathbf{u}^a, \quad (5)$$

where \mathbf{u}^a is the displacement at node \mathbf{X}^a and N^a its associated shape function.

The FE formulation for (4) consists in solving for the array of nodal displacement $\mathbf{u}_h \in \mathbb{R}^3$

$$\mathbf{M} \cdot \ddot{\mathbf{u}}_h + \mathbf{f}_{int} = \mathbf{f}_{ext} \quad (6)$$

where \mathbf{M} , $\mathbf{f}_{int}(\mathbf{u})$, and \mathbf{f}_{ext} are, respectively, the mass matrix and the vector of internal and external forces vector. These quantities are constructed and assembled from the corresponding element quantities \mathbf{M}^e , \mathbf{f}_{int}^e and \mathbf{f}_{ext}^e , which are the mass matrix, internal and external force vectors, respectively. These elemental tensors are calculated using Gauss quadrature approximation in the isoparametric referential and are given (in indicial form) by:

$$M_{iakb}^e = \sum_q w_q \rho_0 \delta_{ik} N_a^e(\boldsymbol{\xi}_q) N_b^e(\boldsymbol{\xi}_q) J_0(\boldsymbol{\xi}_q), \quad (7)$$

$$f_{int,ia}^e = \sum_q w_q P_{iJ}(\boldsymbol{\xi}_q) N_{a,J}^e(\boldsymbol{\xi}_q) J_0(\boldsymbol{\xi}_q), \quad (8)$$

$$f_{ext,ia}^e = \sum_p w_p t_{0i}(\boldsymbol{\xi}_p) N_a^e(\boldsymbol{\xi}_p) J_0(\boldsymbol{\xi}_p) + \sum_q w_q b_{0i}(\boldsymbol{\xi}_q) N_a^e(\boldsymbol{\xi}_q) J_0(\boldsymbol{\xi}_q), \quad (9)$$

where δ_{ik} is the Kronecker symbol, $\boldsymbol{\xi}_q$ corresponds to the coordinates of Gauss point q in the isoparametric referential, w_q is its corresponding weight, and J_0 is the Jacobian between the isoparametric referential and the reference configuration. The same notations with subscripts p (first term of Equation (9)) are for the surface elements belonging to Γ_{d0} mapped onto the corresponding faces of the volume elements.

As the internal force vector \mathbf{f}_{int}^e is not necessarily linear in \mathbf{u} , a Newton-Raphson procedure can be used in conjunction with a time discretization scheme (see Subsection 2.3). In this case, the elemental stiffness matrix \mathbf{K}^e is needed:

$$K_{iakb}^e = \frac{\partial f_{int,ia}^e}{\partial u_{kb}} = \sum_q w_q \frac{\partial P_{iJ}}{\partial F_{kL}}(\boldsymbol{\xi}_q) N_{a,J}^e(\boldsymbol{\xi}_q) N_{b,L}^e(\boldsymbol{\xi}_q) J_0(\boldsymbol{\xi}_q) \quad (10)$$

where the tangent moduli $\frac{\partial P_{iJ}}{\partial F_{kL}}$ are provided by the constitutive law at each time step.

2.3 Time discretization

The generalized Newmark formulation used here for Equation (6) can be written as [24]

$$\mathbf{M} \cdot \ddot{\mathbf{u}}^{n+1} + \mathbf{f}_{int}^{n+1} = \mathbf{f}_{ext}^{n+1}, \quad (11)$$

$$\mathbf{u}^{n+1} = \mathbf{u}^n + \Delta t \dot{\mathbf{u}}^n + (\Delta t)^2 \left[\left(\frac{1}{2} - \beta \right) \ddot{\mathbf{u}}^n + \beta \ddot{\mathbf{u}}^{n+1} \right], \quad (12)$$

$$\dot{\mathbf{u}}^{n+1} = \dot{\mathbf{u}}^n + \Delta t \left[(1 - \gamma) \ddot{\mathbf{u}}^n + \gamma \ddot{\mathbf{u}}^{n+1} \right], \quad (13)$$

where superscripts “ n ” and “ $n+1$ ” indicate that the variables are evaluated at time t^n and t^{n+1} , respectively (note that the subscript “ h ” has been dropped for simplicity). Parameters β and γ set the characteristics of the Newmark

scheme. Parameter $\beta = 0$ leads to an explicit Newmark scheme, whereas for values $0 < \beta \leq 0.5$ leads to an implicit scheme. In the latter case, the set of equations (11)–(13) are solved for the unknown displacement \mathbf{u}^{n+1} , velocity $\dot{\mathbf{u}}^{n+1}$, and acceleration $\ddot{\mathbf{u}}^{n+1}$ using the iterative Newton-Raphson algorithm.

3 Numerical implementation

The Alya system is a computational mechanics code developed with two main motivations. First, it was designed to run with high efficiency standards in large scale supercomputing facilities. Secondly, various physical problem should be allowed to be solved individually or in a coupled manner, while conserving exceptional scalability and retaining their individual efficiency.

Alya's architecture is modular, grouping the different tasks into *kernel*, *modules* and *services*. The *kernel*, the core of Alya, contains all the facilities required to solve any set of discretized PDEs (*e.g.*, the solver, the I/O, the coupling, the elements database, the geometrical information, etc.).

The physical description of a given problem is provided by its corresponding *module* (*e.g.*, the discretized terms of the PDE, the meaning of the boundary and initial conditions, etc.). Other Alya modules than the one presented here include incompressible or compressible flow, thermal transport, or N-body problem, among others. Finally, the *services* contain the toolboxes providing several independent procedures to be called by *modules* and *kernel* (*e.g.*, parallelization or optimization).

3.1 Alya's solid mechanics module

A large database of element types and integration schemes of different orders is available from previous work in the other modules. Well known constitutive equations for large deformation elasticity constitutive models (*i.e.*, neo-Hookean, Holzapfel-Ogden, etc. as well as bridges with usual commercial FE subroutines such as VUMAT [1]) were also implemented.

Figure 1 shows a flowchart of the *solidz* module. All the geometrical and physical data of the problem are introduced as input files. Once read, Alya initializes the computation within the *solidz* module, either in serial or parallel mode. The parallel service must be specified in the input files.

3.2 Mesh convergence

The following test aims at studying the mesh convergence of the Alya-*solidz* module. To this end, a target solution $\mathbf{u}^{(e)}$ with a given degree of regularity is used. By using a stationary manufactured solution belonging to the finite element space (for linear and bilinear elements), the numerical scheme should provide an exact nodalwise solution, *i.e.*, $\mathbf{u} = \mathbf{u}^{(e)}$, thus a) confirming that the equation is well coded and b) allowing for the proposed mesh convergence

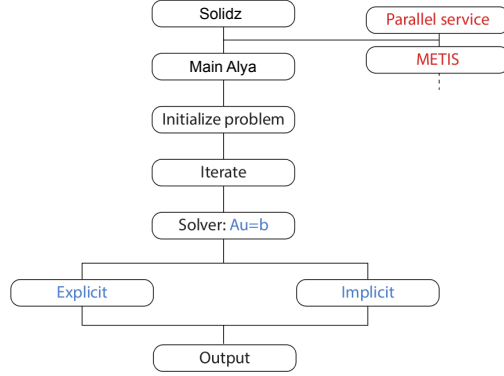


Fig. 1 *solidz* module structure in Alya; *kernel* elements are in blue and *services* in red.

study. The analysis below uses indicial notation and the following equation, obtained from Equation (1), is solved:

$$\rho_0 \ddot{u}_i - \frac{\partial P_{iJ}}{\partial X_J} = - \frac{\partial P_{iJ}^{(e)}}{\partial X_J}, \quad (14)$$

where $P_{ij}^{(e)} = P_{ij}(\mathbf{u}^{(e)})$. Assuming large deformations and using a linear elastic constitutive law, *i.e.*, the last term of Equation (14) is given by:

$$\frac{\partial P_{iJ}^{(e)}}{\partial X_J} = C_{KJML} \left[\frac{F_{iK}^{(e)}}{2} \left(\frac{\partial^2 u_n^{(e)}}{\partial X_J \partial X_M} F_{nL}^{(e)} + \frac{\partial^2 u_n^{(e)}}{\partial X_J \partial X_L} F_{nM}^{(e)} \right) + \frac{\partial^2 u_i^{(e)}}{\partial X_K \partial X_J} E_{ML}^{(e)} \right] \quad (15)$$

where C_{KJML} is the fourth order elasticity tensor and F_{nL} (or F_{iK} and F_{nM}) and E_{ML} are the deformation gradient and the Green-Lagrange strain tensors, defined as:

$$\begin{aligned} F_{mL}^{(e)} &= \frac{\partial u_m^{(e)}}{\partial X_L} + \delta_{mL} \\ E_{ML}^{(e)} &= F_{nM}^{(e)} F_{nL}^{(e)} - \delta_{ML} \end{aligned} \quad (16)$$

The 2D manufactured solution considered here is arbitrarily chosen to be:

$$\begin{cases} u_1^{(e)} = X_1^3 X_2^4 \\ u_2^{(e)} = X_1^3 X_2^3 \end{cases} \quad (17)$$

and is solved on a unit square. Figure 2 shows the resulting mesh convergence results computed on a regular mesh of $Q1$ (quadrilateral) elements. The L^2

and H^1 -norms of the error, $\|\epsilon(\mathbf{u})\|_{L^2}$ and $\|\epsilon(\mathbf{u})\|_{H^1}$, respectively, are shown and they are computed as follows:

$$L^2\text{-norm : } \|\epsilon(\mathbf{u})\|_{L^2} = \sqrt{\int_{\Omega} (\mathbf{u}_h - \mathbf{u}^{(e)})^2 d\Omega} / \sqrt{\int_{\Omega} \mathbf{u}^{(e)2} d\Omega},$$

$$H^1\text{-norm : } \|\epsilon(\mathbf{u})\|_{H^1} = \|\epsilon(\mathbf{u})\|_{L^2} + \|\epsilon(\nabla \mathbf{u})\|_{L^2}.$$




Fig. 2 Mesh Convergence study; L^2 -error and H^1 -error.

The results successfully confirm a first order convergence in H^1 and second order in L^2 with respect to h for the *solidz* module [20].

4 MPI Parallelization

4.1 Parallel context

The parallelization for distributed memory supercomputers is based on a domain decomposition technique, using a Master-Workers paradigm and MPI as the message-passing library. In the core of the parallelization outer layer lies the automatic graph partition tool, METIS [4] (an automatic graph partitioner). First, the master reads the mesh and performs the partition of the mesh into submeshes, or subdomains, using METIS. Each of these subdomains is called a Worker. The Master distributes each partition among the Workers that carry out the solution computational work. METIS mesh partition is done by maximizing load balance and minimizing communication (see Ref. [18] for

more details). As a second step, the Workers build the local elements matrices and local right-hand sides, and are in charge of solving the resulting system solution in parallel. In the assembling tasks, no communication is needed between the Workers, and the scalability only depends on the load balancing. In the iterative solvers, the scalability depends on the size of the interfaces and on the communication scheduling.

Depending on the constitutive model, the resulting equations can be symmetric or non-symmetric. In the implicit case, the basic iterative solvers are GMRes and CG [31]. During the execution of these iterative solvers, two main types of asynchronous communications are required:

- Point-to-point communications via `MPI_IRecv` and `MPI_Isend`, which are used when sparse matrix-vector products are calculated.
- Collective communications via `MPI_Allreduce`, which are used to compute residual norms and scalar products.

In the current implementation of Alya, the solution obtained in parallel is, up to round-off errors, the same as the sequential one all the way through the computation. This is because the mesh partition is only used for distributing work without in any way altering the actual sequential algorithm. This would not be the case if one would consider more complex solvers, like primal or dual Schur complement solvers [7], or more complex preconditioners, like linelet [32] or block LU [26]. Since the explicit framework is relatively straightforward to implement, in the following we only focus our attention on the implicit framework.

The numerical solution of a PDE (and consequently the solid mechanics equations) consists mainly of two steps: first, the construction of the matrix \mathbf{A} and right-hand side (RHS) \mathbf{b} of the algebraic system $\mathbf{Ax} = \mathbf{b}$; second, the solution of this system using an iterative solver. As far as the matrix and RHS assemblies are concerned, only part of the matrix is assembled for the interface nodes. For the iterative solvers, the basic operations are the matrix-vector and the dot products. Let us consider these two operations for a simple one-dimensional case illustrated in Figure 3.

In the context of FE, the coefficients of the matrix come from element computations (see Subsection 2.2). The contribution of node 3 comes from subdomain 1 and 2, namely A_{33}^1 and A_{33}^2 , respectively. Since

$$y_3 = A_{32}x_2 + A_{33}x_3 + A_{34}x_4$$

and by rewriting it as

$$\begin{aligned} y_3 &= (A_{32}x_2 + A_{33}^{(1)}x_3) + (A_{33}^{(2)}x_3 + A_{34}x_4) \\ &= y_3^{(1)} + y_3^{(2)} \end{aligned}$$

the parallelization only consists of a residual exchange.

Therefore, the idea is to use the distributive law of the multiplication to carry out the matrix-vector product in parallel. Let us introduce two functions, that will be described formally in the next subsection as **PAREXC**

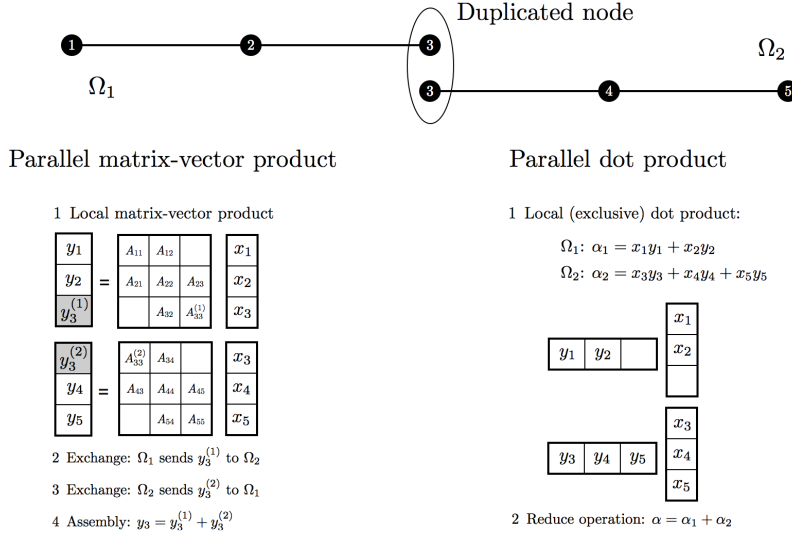


Fig. 3 Iterative solver basic operations: matrix-vector product and dot product.

and **PARASM**. In the previous 1D example, the asynchronous matrix-vector product can be carried out in parallel as follows:

- Perform local matrix-vector product of boundary nodes (node 3);
- **PAREXC**: Exchange the results on the interface $y_3^{(1)}$ and $y_3^{(2)}$ using non-blocking MPI functions;
- Perform local matrix-vector product of interior nodes (rows 1 and 2 for subdomain 1, 4 and 5 for subdomain 2);
- Synchronization `MPI_Waitall`;
- **PARASM**: Assemble (sum) the local contributions of each subdomain: $y_3 = y_3^{(1)} + y_3^{(2)}$.

Regarding the dot product, *i.e.*, $x \cdot y$ we only need to introduce the concept of *own* interface node. In the current implementation all vectors are assembled on the interface, which implies:

$$x_3^{(1)} = x_3^{(2)} = x_3 \text{ and } y_3^{(1)} = y_3^{(2)} = y_3$$

Then, if both subdomains compute their entire local dot product, then the sum of the two contributions will account for $x_3 y_3$ twice, hence

$$\alpha = x_1 y_1 + x_2 y_2 + 2 x_3 y_3 + x_4 y_4 + x_5 y_5,$$

which provides a wrong result. In order to account only once for the interface value, all the interface nodes are splitted into *own* interface nodes and *oth* (for other) interface nodes. That is, an interface node has the *own* status only in one subdomain although it can be shared by others, but with status *oth*. The

local dot product thus only involves interior nodes and interface nodes of *own* status.

In the next section we define formally the concept of interior and interface nodes and introduce some notations and operators, as well as the functions **PAREXC** and **PARASM** doing the parallel matrix-vector and dot products. The term boundary will refer to the interfaces between subdomains.

4.2 Formal set definitions

Before introducing the parallel operators, it is first necessary to define the basic sets, categorizing the nodes and elements of each subdomain, as well as the relations between the subdomains. Some of the definitions are illustrated in Figure 4.




Fig. 4 Node sets.

Consider a computational domain $\Omega \in \mathbb{R}$. Let $\mathcal{N} = \{n_1, n_2, \dots, n_N\}$ and $\mathcal{E} = \{e_1, e_2, \dots, e_E\}$ be the sets of all nodes and elements, respectively, of the FE mesh used to discretize the computational domain. Here, N and E denote the total number of nodes and elements of the computational mesh.

Any element $e \in \mathcal{E}$ can be defined as a tuple of nodes $e = (n_1^e, n_2^e, \dots, n_k^e)$ where k is the number of nodes per element.

Two sets are defined to describe the nodal connectivity with elements and nodes of the mesh, $\mathcal{L}_{ele}(n)$ and $\mathcal{L}_{nod}(n)$, respectively. For any arbitrary node $n \in \mathcal{N}$:

Definition. Element connectivity of n . Let $\mathcal{L}_{ele}(n)$ denote the set of elements in \mathcal{E} directly connected to the node n . Formally,

$$\mathcal{L}_{ele}(n) = \{e \in \mathcal{E} : n \in e\}. \quad (18)$$

Definition. Node connectivity of n . Let $\mathcal{L}_{nod}(n)$ denote the set of nodes in \mathcal{N} directly connected to n . Formally,

$$\mathcal{L}_{nod}(n) = \{m \in \mathcal{N} : \exists e \in \mathcal{L}_{ele}(n), m \in e\} \setminus \{n\} \quad (19)$$

The latter set represents the nodal connection of the mesh and the matrix graph as well. It contains the information required to construct the Compressed Sparse Row (CSR) format used to assemble the sparse matrix of the algebraic system.

We then consider a domain decomposition by elements (see the example of Subsection 4.1). In this work, the METIS [4] library is used. Let \mathcal{N}^I and \mathcal{E}^I denote the set of nodes and elements of any arbitrary subdomain Ω_I , respectively. The total number of nodes and elements of the mesh can be grouped by subdomains:

$$\mathcal{E} = \bigcup_{I=1}^S \mathcal{E}^I, \quad \mathcal{N} = \bigcup_{I=1}^S \mathcal{N}^I,$$

where S is the number of subdomains. As the domain decomposition is performed by elements, the nodes can be shared between subdomains, namely boundary nodes, but the element subdomains are non-overlapping. .

The interface created by the partition of the mesh divides the set of nodes of any arbitrary subdomain Ω_I into two disjoint sets. On the one hand, the set of interior nodes:

Definition. Interior nodes of Ω_I . Let \mathcal{N}_{int}^I denote the set of interior nodes of the subdomain Ω_I . Formally,

$$\mathcal{N}_{int}^I = \mathcal{N}^I \setminus \left(\bigcup_{J=1, J \neq I}^S \mathcal{N}^J \right).$$

These nodes do not belong to the boundary.

On the other hand, the set of boundary nodes:

Definition. Boundary nodes of Ω_I . Let \mathcal{N}_{bou}^I denote the set of boundary nodes of Ω_I . Formally,

$$\mathcal{N}_{bou}^I = \mathcal{N}^I \setminus \mathcal{N}_{int}^I$$

These nodes belong to the interface and are shared by different subdomains, including Ω_I .

In this new context of domain decomposition, let us define a subset of the set $\mathcal{L}_{nod}(n)$ related with any boundary node $n \in \mathcal{N}_{bou}^I$ as:

Definition. Node connectivity of n belonging to Ω_I . Let $\mathcal{L}_{nod}^I(n)$ denote the set of nodes in \mathcal{N}^I directly connected to n . Formally,

$$\mathcal{L}_{nod}^I(n) = \{m \in \mathcal{N}^I : \exists e \in \mathcal{L}_{ele}(n), m \in e\} \setminus \{n\} \quad (20)$$

The concept of adjacency between subdomains is defined as

Definition. Adjacent subdomains. Two arbitrary subdomains I and J are adjacent when $\mathcal{N}^I \cap \mathcal{N}^J \neq \emptyset$.

In order to carry out the scalar product we also need to introduce the concept of *own* boundary and *oth* boundary, where *oth* is defined in Subsection 4.1. This is achieved by splitting the interfaces between the subdomains that share it, for example with METIS.

Definition. Own and other's boundaries of Ω_I . Let us define $\mathcal{N}_{bou,own}^I$ and $\mathcal{N}_{bou,oth}^I$ such that

$$\begin{aligned} \mathcal{N}_{bou,own}^I \cap \mathcal{N}_{bou,oth}^I &= \emptyset, \\ \mathcal{N}_{bou,own}^I \cup \mathcal{N}_{bou,oth}^I &= \mathcal{N}_{bou}^I, \\ \mathcal{N}_{bou,own}^I \cap \bigcup_{J=1, J \neq I}^S \mathcal{N}_{bou,own}^J &= \emptyset. \end{aligned}$$

4.3 Parallel operators

We now define the parallel operator used in the algebraic solver to carry out the matrix-vector product and assembly contributions in an asynchronous way.

The node and element numbering is represented by a local index $index^I$, and an index used to exchange information between two neighboring subdomains I and J , $index_{bou}^{I,J}$.

For any arbitrary subdomain Ω_I ,

Definition. Local node numbering in Ω_I . Let the function

$$index^I : \mathcal{N}^I \rightarrow \{1, 2, 3, \dots, |\mathcal{N}^I|\} \quad (21)$$

be the local node numbering in Ω_I defined as $index^I(n) = i^I$ where $n \in \mathcal{N}^I$ and $i^I \in \mathbb{N}$. $|\mathcal{N}^I|$ is the cardinal number of the set.

For implementation purposes, the interior nodes are first numbered followed by the boundary nodes, as shown in Figure 5. This ordering is useful to carry out not only the matrix-vector product in an efficient way but also the dot product.

Note that for two arbitrary adjacent subdomains Ω_I and Ω_J , $index^I(n)$ is not necessarily equal to $index^J(n)$ for any boundary node $n \in \mathcal{N}_{bou}^I \cap \mathcal{N}_{bou}^J$.




Fig. 5 Node numbering of subdomain Ω_I .

Definition. *Shared boundary node numbering in Ω_I and Ω_J .* Let the function

$$index_{bou}^{I,J} : \mathcal{N}_{bou}^I \cap \mathcal{N}_{bou}^J \rightarrow \{1, 2, 3, \dots, |\mathcal{N}_{bou}^I \cap \mathcal{N}_{bou}^J|\} \quad (22)$$

be the shared boundary node numbering for two arbitrary adjacent subdomains Ω_I and Ω_J defined as $index^{I,J}(n) = i \in \mathbb{N}$ where $n \in \mathcal{N}_{bou}^I \cap \mathcal{N}_{bou}^J$.

This index facilitates the data exchange between two subdomains and is constructed in such a way that both subdomains are able to order their shared boundary nodes in the same way, see Figure 6. In Alya, this task is carried out by the master process, after the partition of the mesh.




Fig. 6 Shared node numbering $index^{I,J}(n)$ of the interface nodes between Ω_I and Ω_J .

Taking into account the previous defined functions, the implementation details of the operators **PAREXC** and **PARASM** from the point of view of any arbitrary subdomain Ω_I are provided in Algorithms 1 and 2. As the exchange of data between subdomains is asynchronous, the operator is divided into two parts: the data exchange first, given by Algorithm 1, and the assembly operation, given by Algorithm 2.

Algorithm 1 The parallel operator **PAREXC** from Ω_I : data exchange

```

Input: a numeric array data with length  $\mathcal{N}^I$ 
Output: void
for each adjacent subdomain  $J$  of  $I$  do
  for each node  $n \in \mathcal{N}_{bou}^I \cap \mathcal{N}_{bou}^J$  do
     $i = index^I(n)$ 
     $i_{bou} = index^{I,J}(n)$ 
    Construct the array  $data\_send(i_{bou}) = data(i)$ 
    MPI_Isend to send the array  $data\_send$  to  $J$ 
    MPI_Irecv to receive the array  $data\_receive^J$  from  $J$ 
  end for
end for

```

Algorithm 2 The parallel operator **PARASM** from Ω_I : data assembly

```

Input: a numeric array data with length  $\mathcal{N}^I$ 
Output: a modified array data
for each adjacent subdomain  $J$  of  $I$  do
  for each node  $n \in \mathcal{N}_{bou}^I \cap \mathcal{N}_{bou}^J$  do
     $i = index^I(n)$ 
     $i_{bou} = index^{I,J}(n)$ 
     $data(i) = data(i) + data\_receive^J(i_{bou})$ 
  end for
end for

```

4.4 Matrix-vector and dot products

The two main parallel functions of an iterative solver, *i.e.*, the matrix-vector and dot products, referred to as **MATVEC** and **DOTPRO** respectively, using the previously defined parallel operators **PAREXC** and **PARASM** are described here. The matrix-vector product is given by Algorithm 3. For the sake of clarity, the indices of the matrix are not given in the CSR format.

Algorithm 4 describes the dot product. As mentioned previously, the loop over the nodes excludes those belonging to the set $\mathcal{N}_{bou,own}^I$ of each subdomain.

It must be emphasized that the functions **PAREXC** and **PARASM** are used to compute many other arrays during the execution of the code, such as the construction of the mass matrix and the diagonal of the stiffness matrix to construct the inverse diagonal preconditioner.

4.5 Scalability test: shear stress of a milling cutter punch

This example, run in implicit, addresses the capability of Alya to deal with structures composed of millions of elements, while maintaining optimal scalability. The test consists in a linear elastic drill (with arbitrary mechanical properties) grasped in a shank. Gravity force and a punctual force along one of its lips are applied. The computational mesh is a regular mesh of 8.5 million tetrahedra, see Figure 7. The tests have been run on MareNostrum cluster.

Algorithm 3 The parallel matrix-vector product **MATVEC**

```

for each subdomain  $I$  do
  for each  $n \in \mathcal{N}_{bou}^I$  do
    for each  $a \in \mathcal{L}^I(n)$  do
       $i = index^I(a)$ 
      Construct  $y^I(i) = 0$ 
      for each  $b \in \mathcal{L}^I(n)$  do
         $j = index^I(b)$ 
        Construct  $y^I(i) = y^I(i) + A^I(i, j) \times x^I(j)$ 
      end for
    end for
  end for
  Exchange: PAREXC( $y^I$ )
end for
for each subdomain  $I$  do
  for each  $n \in \mathcal{N}_{int}^I$  do
    for each  $a \in \mathcal{L}^I(n)$  do
       $i = index^I(a)$ 
      Construct  $y^I(i) = 0$ 
      for each  $b \in \mathcal{L}^I(n)$  do
         $j = index^I(b)$ 
        Construct  $y^I(i) = y^I(i) + A^I(i, j) \times x^I(j)$ 
      end for
    end for
  end for
end for
MPI_Waitall
for each subdomain  $I$  do
  Assemble:  $y^I = \mathbf{PARASM}(y^I)$ 
end for

```

Algorithm 4 The parallel dot product **DOTPRO**

```

for each subdomain  $I$  do
   $\alpha = 0$ 
  for each  $n \in \mathcal{N}^I$  do
     $i = index^I(n)$ 
    if  $n \in \mathcal{N}_{int}^I$  or  $n \in \mathcal{N}_{bou,own}^I$  then
       $\alpha = \alpha + x^I(i) \times y^I(i)$ 
    end if
  end for
end for
MPI_AllReduce of  $\alpha$ 

```

The machine consists of 3,056 nodes, two 8-core processors (Inter Xeon E5-2670 cores at 2.6 GHz) per node and 32 GBytes of memory per node. Figure 8 shows the displacement (left figure) and the maximum and minimum shear stress after 20 time steps. Figure 9 shows the scalability obtained up to 1,024 CPUs. Note that the curve loses its optimallity at 1,024 CPUs, due to the small number of elements per CPU. In the next section the scalability with a finer mesh is studied.




Fig. 7 Scalability test: computational mesh.




Fig. 8 Scalability test: displacement field (left) and maximum and minimum shear stress field (right). The minimum value is shown in blue and the maximum in red.

5 Hybrid based strategy for parallelization

The most important trends in contemporary computer architecture are currently leaning towards processors with multiple cores per socket with access to the same memory bank. This approach allows to run at lower frequencies with better overall performance than a single core processor. The parallelization strategy for this architecture is multithreading: a Master thread forks a number of Worker threads, and the computation is divided among them. The data communication and the synchronization between threads are done using the shared memory inside the multiprocessor. This strategy is known as parallelism at thread level and OpenMP is the standard interface for this model [6].

As detailed in the previous section, parallelization in Alya is mainly done via MPI. The domain decomposition strategy implemented only uses parallelism at task level, which is provided by MPI. For that reason, a hybrid code with OpenMP/MPI is developed in order to take advantage of all levels of parallelism that a multicore architecture offers and also to enable one MPI task to access all the memory of a node. The main feature of this multicore-




Fig. 9 Scalability test: scalability for a fixed number of iterations with a mesh of 8.5 million elements.

architecture relies on the fact that both the shared memory inside the processor and the message passing between nodes are leveraged. The structure consists of two steps: first, it assigns one task to each node and secondly, the node creates a thread per core, see Figure 10. An important reduction of the communication cost between cores is then expected.




Fig. 10 Multicore architecture for an hybrid openMP/MPI framework.

Cost (%)	Routine	Description
...	...	
93,6	iterative scheme	controlling the internal loop of the equations
93,6	explicit time integration	explicit time stepping scheme (Newmark)
93,5	iterative solution	solving an iteration of the equations
93,5	matrix	computing the elemental matrix and RHS
93,5	assembly	assembling the matrix and RHS
36,4	finite element computations	computing derivatives at gauss points
...	...	

Table 1 Routines costs in terms of the sum of time of the function, including the time of the called subfunctions with respect to the total execution time in the explicit case.

5.1 Introducing OpenMP into the MPI parallel code

To exploit the thread-level parallelism of a multicore architecture, OpenMP directives are added to the MPI parallel code. The first step consists in selecting the most-time consuming routines of the code in total execution time. Among these routines, not all of them are susceptible to be parallelized at thread-level: they must contain a loop structure with independent iterations. The targeted loops must contain a large body of code, since a considerable amount of computational work hides behind the overhead of thread-management.

Table 1 depicts the most time-consuming routines of Alya in terms of the total execution time in the explicit case. Note that the cost refers to the percentage of the total execution time spent, also including the call to other subroutines. In order to select the proper subroutine to be optimized by the parallelization with OpenMP, it is important to identify the calls between them. For instance, the routine that computes the matrix costs the same as the assembly one, which means that the time spent by the matrix construction routine itself is minimum, since it has a call to the assembly subroutine. Hence, the assembly routine is the one that has to be optimized. In the FE matrix and RHS assembly routine, the discrete system to be solved is formed by looping over all the elements of the mesh, adding the contributions from that element to the global matrix and RHS. This routine is expensive for a number of reasons: significant loop nesting, many matrices assembling and several calls to other subroutines (for instance, to constitutive models).

In order to introduce thread parallelism in the assembly loop, two main OpenMP directives are used to indicate the compiler how to parallelize appropriately:

- *Guided scheduling*: since the workload within each iteration is not the same, iterations are assigned to threads in blocks. In a guided scheduling the block size decreases within each iteration (in contrast with the dynamic scheduling), thus obtaining a better relation between the thread management time and the balanced workload.
- *Data scope*: the variables that are shared among the iterations are visible to all threads, while the ones that have an independent value among iterations have a different copy per thread.

A critical point on the thread parallelism relies on the so called *race conditions*, which might lead to non-deterministic results. Race conditions arise when different threads try to update the same state or array at the same time. Code paths accessing and manipulating shared data simultaneously are known as *critical regions*. It is clear that both the operation to assemble an element into a local matrix and the addition of that local matrix into the global matrix must be thread safe. Even though these regions cannot be fully avoided it is important to minimize their occurrence, because their abuse can serialize the execution, hence losing the optimal parallel architecture of the code. One possible solution is to specify different region names, combined with the use of *atomic* clauses (uninterruptible instructions) for a single memory location.

Such scheme was implemented in Alya, paying particular attention to the specificities of solid mechanics codes. In order to evaluate the gain in terms of time execution that OpenMP architectures provides, a reduced case was computed. As an illustration of the strategy used here for all identified sub-routines, Figure 11 depicts the speedup analysis of the assembly routine for both explicit and implicit schemes. The performance with one thread is the same as the one for the sequential version and, the thread version for both schemes is speeding up quasi-optimally up to at least 16 cores.




Fig. 11 Scalability of the assembly routine parallelized with OpenMP in Alya-solidz.

6 Numerical examples

This section presents several examples of different nature with the purpose of showing the applicability, scalability, performance and flexibility of Alya when solving a solid mechanics problems.

6.1 Example 1: Mesh Multiplication

In petascale applications, the pre- and post-process tasks are becoming a bottleneck in the complete simulation cycle. Techniques such as parallel I/O have been introduced to mitigate these effects in post-processing, but these are only effective within a limited range. Mesh Multiplication (MM) was introduced as an alternative [17]. This technique consists in refining the mesh uniformly, recursively, on-the-fly and in parallel. For tetrahedra, hexahedra and prisms, each level multiplies the number of elements by eight, while a pyramid is divided into ten new elements. Apart from generating a fine mesh in parallel, the MM strategy enables to carry out the simulation on this fine mesh without having to permanently store it at anytime during the simulation process. This technique is also very useful for studying mesh convergence as well as weak or strong scalability [17]. Figure 12 shows the recursive MM algorithm.




Fig. 12 Example 1: Mesh multiplication algorithm.

The example consists of a linear elastic large structure with arbitrary mechanical properties under its own weight and magnetic load. The initial mesh is composed of a mesh of 491,415 tetrahedral elements. As an example of the efficiency of the algorithm, a mesh of 250 million elements was also obtained in just 1 second on 8,000 CPUs. Figure 13 shows the original mesh and the obtained after MM. Figure 14 depicts the displacement field after 200 iterations using an implicit scheme and the hybrid version of Alya.

Figure 15 shows the speedup obtained with a mesh of 32 million elements (2-level mesh) from 64 to 2,048 CPUs. The code shows optimal scalability when solving mechanical problems. Note that in contrast with the drill problem, the amount of elements per core is significantly larger.




Fig. 13 Example 1: Initial mesh of the structure (left figure) and detail of the mesh after mesh multiplication (right figure). Fusion reactor vacuum vessel. Geometry and mesh property of F4E.

6.2 Example 2: The Chimera method. Two-material cube

As an illustration of the flexibility of the *Alya-solidz* module, the Chimera method is applied to the solid mechanics equations described in Subsection 2.1. This family of methods allows to handle non-conforming and overlapping meshes, thus simplifying the construction of computational meshes of complex geometries.

The origin of the Chimera Method is found in the context of Computational Fluid Dynamics. It was originally developed by Steger and coworkers [33, 11, 34]. It consists in superimposing an independent mesh referred to as the *patch mesh* on a larger mesh covering the overall computational domain and called the *background mesh*. In Alya, the patch mesh and its corresponding constitutive properties are responsible for the mechanical deformation of the body it covers, while the mechanics of the rest of the body is defined by the part of the background mesh not covered by the patch mesh. The coupling between




Fig. 14 Example 1: Displacement field after 200 iterations using and implicit scheme.




Fig. 15 Example 1: Scalability of the fusion reactor problem with a mesh of 32 million elements.

the boundary of the patch mesh and the background mesh is then enforced by additional *transmission conditions*. In the framework of solid mechanics, this non-local approach allows for the independent meshing of complex intertwined geometries without being constrained by mesh conformity conditions at their boundaries. See Ref. [13] for more details.

The Chimera method is applied here to the solution of a 3D example shown in Figure 16. The geometry is composed of two different Neo-Hookean materials, one of which corresponds to the spherical patch mesh enclosed in a

	Elements Total	Elements Extension	Elements Hole	Elements Size
Mesh 1	374,848	14,848	17,120	0.0029
Mesh 2	2,939,664	59,664	151,448	0.00036
Reference mesh	10,790,400			0.000093

Table 2 Example 2: Number of elements and mesh size.

second material. In this example arbitrary constitutive parameters have been considered for both material, but material 2 (within the sphere) is stiffer than material 1. The cube is submitted to an increasing y -displacement on the top face, the bottom face being constrained in the y -direction. Additional lateral boundary conditions are applied to avoid rigid body motion (not shown for clarity).

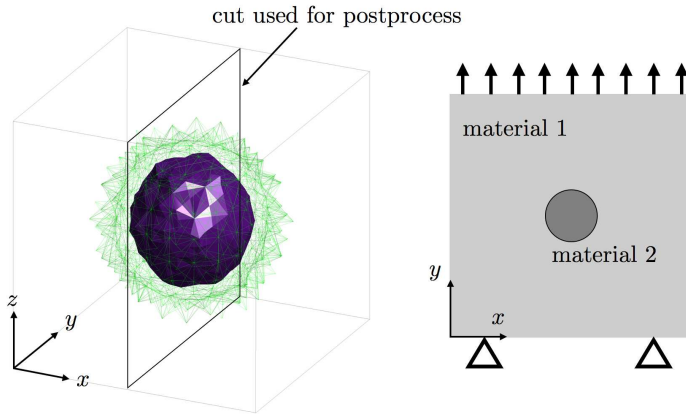


Fig. 16 Example 2: extension elements and hole (left) and boundary conditions (right).

The problem was solved on two different meshes with the Chimera method: a coarse one of 400,000 elements and a fine one of 3 million elements. A reference solution without the Chimera method was also computed in a refined mesh of 10 million elements. Table 2 shows the geometrical details of the three different meshes. Figure 17 shows the displacement and stress in the y direction as well as the corresponding mesh obtained on the mid yz -plane for the coarse and fine meshes of the Chimera method. For the Chimera meshes, the extension elements used to couple the patch and background mesh are delimited in red. Figure 18 shows the solution for an additional cut along the vertical displacement and stress. The solution on both meshes with the Chimera method is compared with the reference one. A good qualitative agreement is observed for the displacement even for the larger mesh. The stress exhibits a lower convergence, but this should be tempered by the fact that the meshes are only

first order for the solution derivatives. In the overlapping zone, covered by the extension elements, two solutions coexist: the one obtained on the patch and the one obtained on the background. The discontinuity in the stress inside the overlapping zone is appreciable for the coarsest mesh, but this jump decreases when refining the mesh. Nevertheless, the maximum error between the solution with Chimera computed with the fine mesh and the reference solution is of order less than 10^{-4} .

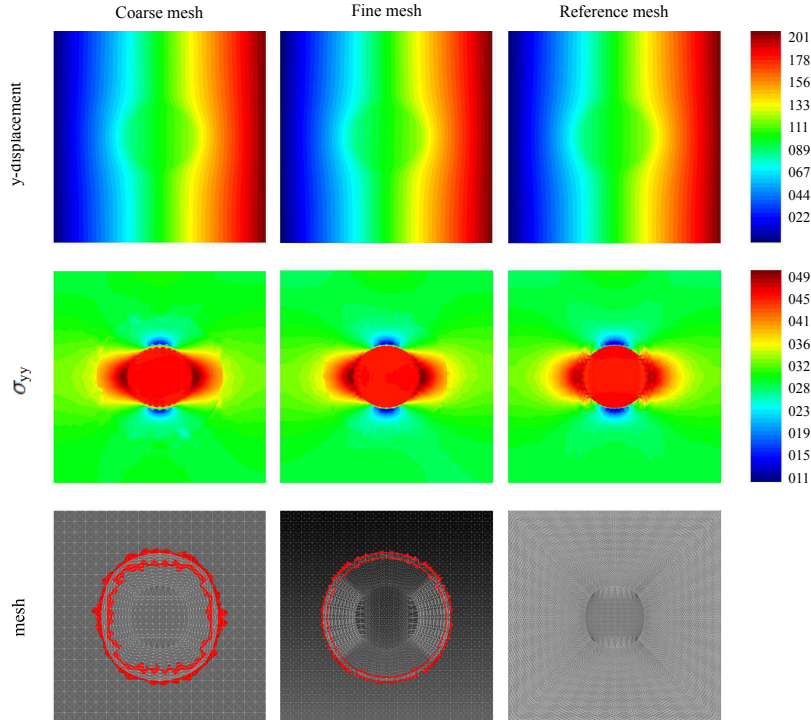


Fig. 17 Example 2: solution on mid yz-plane for the three meshes: y-displacement (top) and stress along the y-direction (middle) and computational mesh (bottom).

6.3 Example 3: Coupled electro-mechanical model of the heart

In this case, Alya is used to simulate the cardiac ventricular contraction, showing the potential for simulating coupled problems. See Ref. [36,23] for a complete discussion of the methodology. The heart is made of elongated cells called myocytes arranged as a compact fibered helicoidal structure. As an electrical impulse propagates, the fibers contract longitudinally making the heart pump the blood out of its cavities thanks to the fiber arrangement. At organ level,




Fig. 18 Example 2: solution along mid line of mid yz-plane for the three meshes: y-displacement (left) and stress along the y-direction (right).

tissue is considered as a continuum composite material with anisotropic behaviour.

The cardiac computational model can be decomposed into three parts:

- **Electrophysiology:** Action potential propagation $\phi(x_i, t)$ is modelled through a transient anisotropic diffusion equation with a non-linear reaction term:

$$\frac{\partial \phi}{\partial t} = \frac{\partial}{\partial x_i} \left(D_{ij} \frac{\partial \phi}{\partial x_j} \right) + L(\phi) \quad (23)$$

Diffusion is governed by the tensor D_{ij} , which represents the fiber orientation at each physical point. Its diagonal components are the axial and crosswise fibre diffusion. The crosswise diffusion is around one third of the axial diffusion. Finally, $L(\phi_\alpha)$ is the non-linear reaction term corresponding to the cell model. Depending on the model used, this term ranges from a simple cubic equation to an ordinary differential equation system of one hundred equations coupled and solved simultaneously. This term models the ionic currents interaction behind the muscular cellular mechanisms.

- **Mechanical deformation:** From a mechanical point of view, cardiac tissue is a thick layered structure: endo-, epi- and myocardium. The material is compressible hyperelastic, with anisotropic behaviour ruled by the fibre structure. The composite character of the material is determined by the internal stress, which is developed in two parts, active and passive:

$$\boldsymbol{\sigma} = \boldsymbol{\sigma}_{pas} + \sigma_{act}(\lambda, [Ca^{2+}]) \mathbf{f} \otimes \mathbf{f} \quad (24)$$

The passive part is governed by a transverse isotropic exponential strain energy function $W(b)$ that relates the Cauchy stress $\boldsymbol{\sigma}$ to the right Cauchy-Green deformation \mathbf{b} [23]:

$$J\boldsymbol{\sigma}_{pas} = (a e^{b(I_1-3)} - a)\mathbf{b} + 2a_f(I_4 - 1)e^{b_f(I_4-1)^2} \mathbf{f} \otimes \mathbf{f} + K(J - 1)\mathbf{I} \quad (25)$$

The strain invariant I_1 represents the non-collagenous material while strain invariant I_4 represents the stiffness of the muscle fibers, and a, b, a_f, b_f are parameters to be determined experimentally. K is the bulk modulus and

\mathbf{f} defines the fibre direction. The active part comes through the electro-mechanical coupling and is described as follows.

- **Electro-mechanical Coupling:** The contracting electrical component of the electro-mechanical coupling is initiated almost simultaneously in several zones of the ventricular epicardium. Cardiac mechanical deformation is the result of the active tension generated by the myocytes. The model assumes that the active stress is produced only in the direction of the fibre and depends on the calcium concentration of the cardiac cell:

$$\sigma_{act} = \gamma \frac{[Ca^{2+}]^n}{[Ca^{2+}]^n + C_{50}^n} \sigma_{max} (1 + \beta(\lambda_f - 1)). \quad (26)$$

In this equation, C_{50}^n , σ_{max} , λ_f and $0 < \gamma < 1$ are model parameters.

Both electrophysiology and mechanical action is simulated in Alya on the same mesh, as fine as the case demands. The time integration scheme is a staggered, with both problems solved explicitly. The fibre fields can come either from measurements (using a so-called Diffusion Tensor MRI technique) or from semi-empirical rule-based models [23].

Figure 19 shows several snapshots of a bi-ventricular geometry during systole. Figure 20 shows the evolution of the ejection rate, which represents the heart pumping action, measuring the change in the ventricular cavities volume.




Fig. 19 Example 3: Bi-ventricular electrical activation propagation of a dog heart during systole.




Fig. 20 Example 3: Heart ejection rate.

6.4 Example 4: Fluid-structure interactions. Wall's problem

This problem [38] is proposed to demonstrate the ability to solve FSI problems with complex flow-flexible structure interactions exhibiting large deformations.

The problem consists of a thin elastic nonlinear shell attached to a fixed square rigid body, which are submerged in an incompressible fluid flow. Both media have arbitrary mechanical properties but such that the Reynolds number is $Re = 204$ if the length of the square rigid body is taken as the characteristic length. Inflow boundary conditions are imposed on the left wall of the fluid domain, outflow on the right wall and slip boundary conditions at the top and bottom wall, see Figure 21. Non-slip boundary conditions are imposed along the body and the structure. Again, both problems are solved using Alya.




Fig. 21 Example 4: Boundary conditions.

The finite element mesh is shown in Figure 22. It consists of a hybrid mesh of 6,919 elements, composed of tetrahedra and quadrilateral elements. The fluid domain contains 6,695 elements, while the solid one has 224 elements.

A weak coupling scheme [35, 8] was used to couple the fluid with the structure. The key point of the strategy implemented in Alya is to divide geometrically the problem in non-overlapping fluid and solid zones, such that by




Fig. 22 Example 4: Reference finite element mesh.

properly distributing the amount of processors working within each zone, the amount of computations performed by the CPUs is optimized.

Figure 23 shows the structural displacement at different times $t = 0, 3.58, 7.01$ and $15.58s$ along with the fluid velocity. Note that the mesh is adapted to the solid structure displacement, since the ALE formulation used here computes the convective velocity of the fluid according to the difference between the material and mesh velocities. The results can be compared to the ones obtained in Ref. [38].




Fig. 23 Example 4: Fluid velocity field for different times.

7 Conclusion

This paper presents the solid mechanics module of Alya code for solving linear and non-linear continua problems with large deformations addressing also the potential for solving complex coupled problems in parallel. The parallelization strategy used in Alya is formally described in the context of MPI communications. To enhance the massively parallel performance of the code, OpenMP directives were added to the current MPI Alya code. The code was tested successfully using thousands of CPUs to solve problems with up to millions of elements. The flexibility of solid module in Alya was demonstrated by using more complex preprocessing techniques, such as a mesh multiplication algorithm and the Chimera method. Those were successfully tested with three-dimensional structural tests showing optimal results. Examples of FSI and electro-mechanical coupling were also shown, demonstrating the capability of the code for solving multi-physic problems with high-performance computing. Future work is aiming at implementing more complex FE techniques and overcoming the barrier of more than 100,000 CPUs.

Acknowledgements D.D.T and A.J acknowledge funding through SIMUCOMP and ERA-NET MATERA+ project financed the Consejería de Educación y Empleo of the Comunidad de Madrid and by the European Union's Seventh Framework Programme (FP7/2007-2013). This work was partially supported by the grant SEV-2011-00067, Severo Ochoa Program, awarded by the Spanish Government.

References

1. Abaqus. URL <http://www.3ds.com/products-services/simulia/portfolio/abaqus/>
2. Alya system. URL <http://www.bsc.es/computer-applications/alya-system>
3. CODE SATURN. URL <http://code-saturne.org>
4. Metis, family of multilevel partitioning algorithms. URL <http://glaros.dtc.umn.edu/gkhome/views/metis>
5. Openfoam. URL <http://www.openfoam.com/>
6. openmp. URL <http://openmp.org/wp/>
7. Absorbing interface conditions for domain decomposition methods: A general presentation. *Computer Methods in Applied Mechanics and Engineering* **195**(29-32), 3880 – 3900 (2006)
8. A.G. Malan, O.O.: An accelerated, fully-coupled, parallel 3d hybrid finite-volume fluid-structure interaction scheme. *Comput. Methods Appl. Mech. Engrg* **253**, 426–438 (2013)
9. Arbenz, P., van Lenthe, G., Mennel, U., Müller, R., Sala, M.: A scalable multi-level preconditioner for matrix-free μ -finite element analysis of human bone structures. *International Journal for Numerical Methods in Engineering* **73**, 937–947 (2008). DOI 10.1002/nme.2101
10. Becker, G., Noels, L.: A full-discontinuous galerkin formulation of nonlinear kirchhoff-love shells: elasto-plastic finite deformations, parallel computation, and fracture applications. *International Journal for Numerical Methods in Engineering* **93**, 80–117 (2013). DOI 10.1002/nme.4381
11. Benek, J.: Chimera. a grid-embedding technique. Tech. rep., DTIC Document (1986)
12. Cirak, F., Deiterding, R., Mauch, S.: Large-scale fluidstructure interaction simulation of viscoplastic and fracturing thin-shells subjected to shocks and detonations. *Computers and Structures* **85**, 1049–1065 (2007). DOI 10.1016/j.compstruc.2006.11.014

13. Eguzkitza, B., Houzeaux, G., Aubry, R., Owen, H., Vázquez, M.: A parallel coupling strategy for the chimera and domain decomposition methods in computational mechanics. *Computers & Fluids* (2013)
14. Flaig, C., Arbenz, P.: A scalable memory efficient multigrid solver for micro-finite element analyses based on CT images. *Parallel Computing* **37**, 846–854 (2011). DOI 10.1016/j.parco.2011.08.001
15. Gerstenberger, A., Tuminaro, R.: An algebraic multigrid approach to solve extended finite element method based fracture problems. *International Journal for Numerical Methods in Engineering* **94**, 248–272 (2013). DOI 10.1002/nme.4442
16. Goudreau, G., Hallquist, J.: Recent developments in large-scale finite element Lagrangian hydrocode technology. *Computer Methods in Applied Mechanics and Engineering* **33**, 725–757 (1982). DOI 10.1016/0045-7825(82)90129-3
17. Houzeaux, G., de la Cruz, R., Owen, H., Vázquez, M.: Parallel uniform mesh multiplication applied to a navier-stokes solver. *Computers & Fluids* **80**, 142–151 (2013)
18. Houzeaux, G., Vázquez, M., Aubry, R., Cela, J.: A massively parallel fractional step solver for incompressible flows. *JCP* **228**(17), 6316–6332 (2009)
19. Hughes, T., Ferencz, R.: Large-scale vectorized implicit calculations in solid mechanics on a CrayX-MP/48 utilizing EBE preconditioned conjugate gradients. *Computer Methods in Applied Mechanics and Engineering* **61**, 215–248 (1987). DOI 10.1016/0045-7825(87)90005-3
20. Hughes, T.J.: The finite element method: linear static and dynamic finite element analysis. Dover Publications. com (2012)
21. Hussain, M., Abid, M., Ahmad, M., Khokhar, A., Masud, A.: A parallel implementation of ALE moving mesh technique for FSI problems using OpenMP. *International Journal of Parallel Programming* **39**, 717–745 (2011). DOI 10.1007/s10766-011-0168-3
22. Komatitsch, D., Erlebacher, G., Göddeke, D., Michéa, D.: High-order finite-element seismic wave propagation modeling with MPI on a large GPU cluster. *Journal of Computational Physics* **229**, 7692–7714 (2010). DOI 10.1016/j.jcp.2010.06.024
23. Lafortune, P., Arís, R., Vázquez, M., Houzeaux, G.: Coupled electromechanical model of the heart: Parallel finite element formulation. *International Journal for Numerical Methods in Biomedical Engineering* **28**, 72–86 (2012). DOI 10.1002/cnm.1494
24. and W.K. Liu, T.B., Moran, B.: *Nonlinear Finite Elements for Continua and Structures*. Wiley (2000)
25. Liu, Y., Zhou, W., Yang, Q.: A distributed memory parallel element-by-element scheme based on Jacobi-conditioned conjugate gradient for 3D finite element analysis. *Finite Elements in Analysis and Design* **43**, 494–503 (2007). DOI 10.1016/j.finel.2006.12.007
26. Maurer, D., Wieners, C.: A parallel block {LU} decomposition method for distributed finite element matrices. *Parallel Computing* **37**(12), 742 – 758 (2011). DOI <http://dx.doi.org/10.1016/j.parco.2011.05.007>
27. Moore, D., Jérusalem, A., Nyein, M., Noels, L., Jaffee, M., Radovitzky, R.: Computational biology — modeling of primary blast effects on the central nervous system. *NeuroImage* **47**(2), T10–T20 (2009). DOI 10.1016/j.neuroimage.2009.02.019
28. R. Quey, P.D., Barbe, F.: Large-scale 3D random polycrystals for the finite element method: Generation, meshing and remeshing. *Computer Methods in Applied Mechanics and Engineering* **200**, 1729–1745 (2011). DOI 10.1016/j.cma.2011.01.002
29. Radovitzky, R., Seagraves, A., Tupek, M., Noels, L.: A scalable 3D fracture and fragmentation algorithm based on a hybrid, discontinuous galerkin, cohesive element method. *Computer Methods in Applied Mechanics and Engineering* **200**, 326–344 (2011). DOI 10.1016/j.cma.2010.08.014
30. van Rietbergen, B., Weinans, H., Huiskes, R., Polman, B.: Computational strategies for iterative solutions of large FEM applications employing voxel data. *International Journal for Numerical Methods in Engineering* **39**, 2743–2767 (1996). DOI 10.1002/(SICI)1097-0207(19960830)39:16<2743::AID-NME974>3.3.CO;2-1
31. Saad, Y.: *Iterative Methods for Sparse Linear Systems*. Society for Industrial and Applied Mathematics (2003)
32. Soto, O., Löhner, R., Camelli, F.: A linelet preconditioner for incompressible flow solvers. *Int. J. Num. Meth. Heat Fluid Flow* **13**(1), 133–147 (2003)
33. Steger, J., Benek, F.D.J.: A chimera grid scheme. *Advances in Grid Generation* **5**, 59–69 (1983)

34. Steger, J., Benek, J.: On the use of composite grid schemes in computational aerodynamics. *Comp. Meth. Appl. Mech. Eng.* **64**, 301–320 (1987)
35. V. Kalro, T.E.T.: A parallel 3d computational method for fluid-structure interactions in parachute systems. *Comput. Methods Appl. Mech. Engrg* **190**, 1467–1482 (2000)
36. Vázquez, M., Arís, R., Houzeaux, G., Aubry, R., Villar, P., Garcia-Barnós, J., Gil, D., Carreras, F.: A massively parallel computational electrophysiology model of the heart. *International Journal for Numerical Methods in Biomedical Engineering* **27**(12), 1911–1929 (2011). DOI 10.1002/cnm.1443
37. Vázquez, M., Houzeaux, G., Grima, R., Cela, J.: Applications of parallel computational fluid mechanics in MareNostrum supercomputer: Low-mach compressible flows. In: *PARCFD2007. Antalya (Turkey)* (2007)
38. Wall, W.A., Ramm, E.: Fluid-structure interaction based upon stabilized (ale) finite element method. In: *IV World Congress on Computational Mechanics. Barcelona. CIMNE.*