

## HPC compact quasi-Newton algorithm for interface problems

Alfonso Santiago<sup>a</sup>, Miguel Zavala-Aké<sup>c</sup>, Ricard Borell<sup>a</sup>, Guillaume Houzeaux<sup>a</sup>,  
Mariano Vázquez<sup>a,b</sup>

5

<sup>a</sup>*Barcelona Supercomputing Center (BSC), Barcelona, Spain.*

<sup>b</sup>*ELEM Biotech, Barcelona, Spain*

<sup>c</sup>*Kungliga Tekniska Högskolan (KTH), Stockholm, Sweden.*

---

### Abstract

In this work we present a robust interface coupling algorithm called Compact Interface quasi-Newton (CIQN). It is designed for computationally intensive applications using an MPI multi-code partitioned scheme. The algorithm allows to reuse information from previous time steps, feature that has been previously proposed to accelerate convergence. Through algebraic manipulation, an efficient usage of the computational resources is achieved by: avoiding construction of dense matrices and reduce every multiplication to a matrix-vector product and reusing the computationally expensive loops. This leads to a compact version of the original quasi-Newton algorithm. Altogether with an efficient communication, in this paper we show an efficient scalability up to 4800 cores. Three examples with qualitatively different dynamics are shown to prove that the algorithm can efficiently deal with added mass instability and two-field coupled problems. We also show how reusing histories and filtering does not necessarily makes a more robust scheme and, finally, we prove the necessity of this HPC version of the algorithm. The novelty of this article lies in the HPC focused implementation of the algorithm, detailing how to fuse and combine the composing blocks to obtain a scalable MPI implementation. Such an implementation is mandatory in large scale cases, for which the contact surface cannot be stored in a single computational node, or the number of contact nodes is not negligible compared with the size of the domain. 2020 Elsevier. This manuscript version is made available under the CC-BY-NC-ND 4.0 license

<http://creativecommons.org/licenses/by-nc-nd/4.0/>

*Keywords:* coupling scheme, fluid-structure interaction, high performance

10 computing, partitioned scheme

*2010 MSC:* 68U20, 00A72, 68Q85, 65-04, 74F10

---

## 1. Introduction

Interface problems like, solid-solid contact, fluid-structure interaction (FSI) or heat transfer gained great attention in the last decades due to the broad range of applications in aerospace industry, manufacturing, wind energy production or biomechanics. These problems can be mathematically treated using heterogeneous domain decomposition methods [1]. In each subdomain the problems are defined with their own Neumann and Dirichlet boundary conditions, including the boundary condition at the contact surface. From the algorithmic point of view, the problem can be attacked with the monolithic or the partitioned scheme. On the former, using an *ad-hoc* solver, one matrix is built including the degrees of freedom for both the fluid and solid [2, 3, 4]. On the latter, fluid and the solid are computed independently as black-box solvers, exchanging the quantities of interest at certain synchronisation points of the workflow [5, 6, 7, 8]. Both strategies have advantages and drawbacks. Monolithic schemes have less numerical instabilities, but leads to a linear system hard to preconditionate and require to design a specific solver from scratch for every pair of coupled problems [2, 4, 9, 10]. The partitioned scheme allows code reusing, but require convergence iterations at each time step [5, 6, 7, 8].

The algorithm presented in this work is based on the interface quasi-Newton method (IQN) [5], improved and extended with a special care in the parallel implementation. The proposed scheme is implemented in Alya [11, 12, 13], the BSC's in-house tool for multiphysics problems. The uncoupled physics solvers in Alya have an almost linear scalability proven up to a hundred thousand cores [14]. Following a black-box multi-code strategy, the coupled problems are solved by executing two different MPI-based parallel instances of Alya which

interchange data in the contact surface.

The goal of this work is to develop both an accurate and efficient version of the interface quasi Newton algorithm. There are three main motivations to do so. Firstly, we look for a coupling strategy that can deal with large-scale problems on both sides of the coupled problem. Secondly, we need a coupling algorithm that is able to robustly tackle the main issue on the fluid-structure interaction (FSI) problem, namely added mass instability [15, 16]. And, finally, the algorithm should be able to deal with more than one coupled interface at the same time (n-field coupling) [17]. When modelling biomechanics, the last two conditions are mandatory, as the densities of the tissues and fluids are similar and multiple cavities are interconnected. A parallel formulation of the coupling algorithm has not been shown in the past, but it is a requirement for massively parallel applications so it does not become the bottleneck. Previous publications [5, 6, 7, 8, 18, 19] assume a negligible cost of the coupling algorithm. This is based on the fact that in some cases the coupled surface might be small compared with the solved volumes. This assumption eases the software development as the code can be written as serial. This hypothesis might be true in some cases, but not in biomechanical applications where the geometrical complexity of the biological structures enormously increases the area of the contact surface. Although it is true that the convergence acceleration requires less operations than the solvers in the Piccard iteration, the interface quasi-Newton algorithm requires a large number of matrix-matrix multiplications. This is intractable for large problems that must run in distributed memory systems where the interface cannot be stored in a single shared-memory node, so an efficient parallelisation is mandatory.

This work is organised as follows. Section 2 contains the mathematical development for the algorithm, detailing the included improvements and the parallelisation strategy. Section 3 shows two experiments with remarkably different dynamics and a scalability test. Conclusions can be found in section 4.

## 2. Material and Methods

In this section, a compact version of the Interface quasi-Newton algorithm is presented. Section 2.2 shows the original algorithm, and section 2.3 the included improvements. The QR decomposition, a critical step in the algorithm, is thoroughly detailed in section 2.4 and the parallelisation strategy in section 2.5. An explanation of the used Einstein index notation convention can be found in Appendix B.

### 2.1. Problem setting

In this work, we focus in surface problems that can be stated as  $d_\alpha^{I+1} = S(f_\alpha)$  and  $f_\alpha^{I+1} = F(d_\alpha)$ . Each form represents the numerical result of a physical problem and  $d_\alpha$  and  $f_\alpha$  are the unknowns at the interface  $\Gamma_c$ . This can also be written as the fixed point equation  $d_\alpha^{I+1} = S(F(d_\alpha))$ , or in a generic manner:

$$\tilde{x}_\alpha^{I+1} = H(x_\alpha), \quad (1)$$

where  $H(x_\alpha)$  condenses both solvers. The parallel solvers  $S(f_\alpha)$  and  $F(d_\alpha)$  can be executed either one after the other in a block-sequential manner (Gauss-Seidel) or at the same time in a block-parallel manner (Jacobi) [20]. While the former is less computationally efficient, it improves convergence of the iterative solver, and therefore will be the used scheme. Performance can be improved with a convergence acceleration algorithm. Examples of them can be found in [21]. In the following section we will develop a high-performance version of an interface quasi-Newton algorithm.

### 2.2. General Overview of the Algorithm

The first implementation of the Interface Quasi Newton (IQN) algorithm is described in [5]. Distinctly to other quasi-Newton schemes, in the IQN the Jacobian is approximated by a field defined in the contact surface and depending on the local residual variation over a given number of iterations [22]. The residual of eq. (1) can be defined as  $r_\alpha = H(x_\alpha) - x_\alpha = \tilde{x}_\alpha - x_\alpha$ . For each time

step, the problem is converged when  $r_\alpha = 0$ . If the Jacobian  $\partial r_\alpha / \partial x_\beta$  is known, the increment of the variable  $x_\alpha$  can be computed as:

$$\frac{\partial r_\alpha}{\partial x_\beta} \Delta x_\beta = -r_\alpha. \quad (2)$$

Therefore, computing the next iterate as  $x_\alpha^{I+1} = x_\alpha + \Delta x_\alpha$ . Generally the exact Jacobian cannot be computed or it is computationally expensive to do so. This quasi-Newton scheme provides a method to obtain an approximation of the inverse Jacobian. The multi-secant equation for the inverse Jacobian reads:

$$\left( \frac{\partial r_\alpha}{\partial x_\beta} \right)^{-1} V_{\alpha i} \approx W_{\beta i}, \quad (3)$$

where:

$$V_{\alpha i} = [\Delta r_\alpha^1, \Delta r_\alpha^2, \dots, \Delta r_\alpha^q] \quad \text{with} \quad \Delta r_\alpha^I = r_\alpha^{I+1} - r_\alpha^I \quad (4)$$

$$W_{\alpha i} = [\Delta \tilde{x}_\alpha^1, \Delta \tilde{x}_\alpha^2, \dots, \Delta \tilde{x}_\alpha^q] \quad \text{with} \quad \Delta \tilde{x}_\alpha^I = \tilde{x}_\alpha^{I+1} - \tilde{x}_\alpha^I. \quad (5)$$

where  $r_\alpha^{I+1}$  and  $\tilde{x}_\alpha^{I+1}$  and  $r_\alpha^I$  and  $\tilde{x}_\alpha^I$  are the current and past values respectively.  $V_{\alpha i}, W_{\alpha i} \in \mathbb{R}^{p \times q}$ , where  $p$  is the number of contact degrees of freedom and  $q$  is the number of saved non-zero iterations where, generally,  $p \gg q$ . Note that the newest values are stored at the left side of the built matrix, while the older values are moved to the right. The residual increment of the current iteration is approximated as a linear combination of the previous residuals increments:

$$\Delta r_\alpha = V_{\alpha i} \lambda_i, \quad (6)$$

where  $\lambda_i \in \mathbb{R}^{q \times 1}$  is the solution of the optimisation problem  $\| \Delta r_\alpha - V_{\alpha i} \lambda_i \|$  described in [23]. To obtain  $\lambda_i$ , the matrix  $V_{\alpha i}$  is decomposed in an orthogonal matrix  $Q_{\alpha\beta} \in \mathbb{R}^{p \times p}$  and an upper triangular  $U_{\alpha i} \in \mathbb{R}^{p \times p}$  with a QR decomposition:

$$V_{\alpha i} = Q_{\alpha\beta} U_{\beta i}. \quad (7)$$

As  $U_{\alpha i}$  is upper triangular, only its first  $q$  rows are different from zero. With this information we can build a modified QR decomposition with  $U_{ij} \in \mathbb{R}^{q \times q}$

and  $Q_{\alpha i} \in \mathbb{R}^{p \times q}$  such that:

$$V_{\alpha i} = Q_{\alpha k} U_{ki}, \quad (8)$$

reducing the amount of memory and computing effort required, as described ahead in section 2.5. After this decomposition, the vector  $\lambda_i$  can be obtained by backsubstitution of the upper triangular matrix  $U_{ij}$ :

$$U_{ij} \lambda_j = Q_{\alpha i} \Delta r_{\alpha}. \quad (9)$$

As  $Q_{\alpha i}$  is orthogonal, the inverse is equal to the transpose, avoiding the inversion of this matrix. Also, as  $\Delta r_{\alpha} = r_{\alpha}^{I-1} - r_{\alpha}$  and the objective is to get  $\Delta r_{\alpha} = 0_{\alpha} - r_{\alpha}$ , we can say:

$$U_{ij} \lambda_j = -Q_{\alpha i} r_{\alpha}. \quad (10)$$

Once  $\lambda_i$  is computed, the increment of the unknown  $\Delta x_{\alpha}$  can be computed as  $\Delta x_{\alpha} = W_{\alpha i} \lambda_i$ , and the update of the unknown as:

$$x_{\alpha}^{I+1} = \tilde{x}_{\alpha} + W_{\alpha i} \lambda_i. \quad (11)$$

The scheme is summarised in algorithm 1. For each time iteration, an initial guess and residue are computed. As the proposed algorithm requires increments, a first step with fixed relaxation  $\omega_0$  is required. After, IQN loop continues until convergence is achieved.

```

1 For each time step, solve:
2  $x_\alpha^0 = x_\alpha^{ini}$ 
3  $\tilde{x}_\alpha^0 = H(x_\alpha^0)$ 
4  $r_\alpha^0 = \tilde{x}_\alpha^0 - x_\alpha^0$ 
5 update  $x_\alpha^1 = \tilde{x}_\alpha^0 + \omega_0 r_\alpha^0$ 
   while problem not converged do
6    $\tilde{x}_\alpha = H(x_\alpha)$ 
7    $r_\alpha = \tilde{x}_\alpha - x_\alpha$ 
8   build  $V_{\alpha i} = [\Delta r_\alpha^{I-1}, \dots, \Delta r_\alpha^0]$ ; with  $\Delta r_\alpha^I = r_\alpha^I - r_\alpha$ 
9   build  $W_{\alpha i} = [\Delta \tilde{x}_\alpha^{I-1}, \dots, \Delta \tilde{x}_\alpha^0]$ ; with  $\Delta \tilde{x}_\alpha^I = \tilde{x}_\alpha^I - \tilde{x}_\alpha$ 
10  decompose  $V_{\alpha j} = Q_{\alpha i} U_{ij}$  (by QR decomposition)
11  solve  $U_{ij} \lambda_j = -Q_{\alpha i} r_\alpha$ 
12  update  $x_\alpha^{I+1} = \tilde{x}_\alpha + W_{\alpha i} \lambda_i$ 
   end

```

**Algorithm 1:** Interface quasi-Newton algorithm overview.

### 2.3. Improvements on the original scheme

It has been proposed [19] that adding information from iterations from the previous time steps into matrices  $V_{\alpha i}$  and  $W_{\alpha i}$  improve the convergence properties of the algorithm. To do so, we redefine matrices 4,5, with information of the iterations from previous time steps:

$$V_{\alpha i} = [V_{\alpha i}^1, V_{\alpha i}^2, \dots, V_{\alpha i}^T,] \quad \text{with} \quad V_{\alpha i}^t \quad \text{as eq. (4)} \quad (12)$$

$$W_{\alpha i} = [W_{\alpha i}^1, W_{\alpha i}^2, \dots, W_{\alpha i}^T,] \quad \text{with} \quad W_{\alpha i}^t \quad \text{as eq. (5)}, \quad (13)$$

where  $t$  ranges from the current processed time step to the last saved time step,  $T$ . Note that  $V_{\alpha i}, W_{\alpha i} \in \mathbb{R}^{p \times q}$ , but now  $q$  is the number of saved non-zero iterations from the current and past time steps. Including this information increases the probability that the columns in  $V_{\alpha i}$  are linearly dependent, rendering the QR decomposition unstable. Different filtering techniques have been proposed

[19] to remove these columns, but all of them require building dense intermediate matrices, or even finding every associated eigenvalue [24], with its associated expensive computational cost. Moreover there is not a clearly better filtering technique [18, 19]. This is why we choose to reuse the simple and parallelised incomplete QR decomposition developed in this work to check the linear dependency of the columns of  $V_{\alpha i}$ . If  $|U_{ii}| < \epsilon \|U\|_2$ , where  $U_{ij}$  is the upper triangular matrix and  $\epsilon$  a parameter, the  $i$ -th column is deleted from  $V_{\alpha i}$  and  $W_{\alpha i}$ . The column deleted might correspond to the current processed time step (sub-matrix  $V_{\alpha i}^1$ ) or any other column corresponding to any other time step (sub-matrix  $V_{\alpha i}^j$ ). This requires re-stacking the non-zero columns to obtain again a dense set of matrices.

#### 2.4. QR decomposition

A critical step is the QR decomposition (line 10 in algorithm 1) due to the numerous matrix-matrix products involved in it. In this section the QR decomposition will be explained and through algebraic manipulation these matrix products will be simplified in the following section. The goal of the QR decomposition is to obtain the orthogonal and the upper triangular matrices  $Q_{\alpha\beta}$  and  $U_{\alpha i}$ , with the following shape:

$$Q_{\alpha\epsilon} = {}^1B_{\alpha\beta} {}^2B_{\beta\gamma} \dots {}^qB_{\gamma\epsilon} \quad (14)$$

$$U_{\alpha i} = {}^qB_{\alpha\beta} \dots {}^2B_{\beta\gamma} {}^1B_{\gamma\epsilon} V_{\epsilon i}, \quad (15)$$

where  $B_{\alpha\beta} \in \mathbb{R}^{p \times p}$  are dense intermediate matrices obtained during the iterative decomposition. At each iteration, the matrix  $V_{\alpha i}$  is processed column by column. We use a left superscript to identify the corresponding iteration of the QR algorithm but, for easiness on the reading, we avoid using any other time or coupling iteration superscripts.  $V_{\alpha i}$  can be considered as a set of  $q$  ordered vectors:

$${}^1V_{\alpha i} = \left[ \begin{array}{c|c|c} \begin{bmatrix} v_{11} \\ v_{21} \\ \vdots \end{bmatrix} & \begin{bmatrix} v_{12} \\ v_{22} \\ \vdots \end{bmatrix} & \cdots & \begin{bmatrix} v_{1q} \\ v_{2q} \\ \vdots \end{bmatrix} \end{array} \right] = [v_{\alpha 1}, v_{\alpha 2}, \dots, v_{\alpha q}]. \quad (16)$$

The algorithm, iteratively makes each column orthogonal to each other column in the matrix. It starts iteration  $j$  with a matrix  ${}^jV_{\alpha i}$  obtained with data from iteration  $j-1$ . To decompose the  $j$ -th column of  ${}^jV_{\alpha i}$ , a unitary vector  $u_\alpha$  has to be built:

$$u_\alpha = \frac{n_\alpha}{\|n_\alpha\|} \quad \text{with,} \quad n_\alpha = v_\alpha - \|v_\alpha\| {}^j e_\alpha, \quad (17)$$

where  $v_\alpha$  is the column to decompose and  ${}^j e_\alpha$  is a unitary vector with  $j$ -th position equal to 1 and to 0 otherwise. Then,

$${}^j B_{\alpha\beta}^* = \delta_{\alpha\beta} - 2u_\alpha u_\beta \quad (18)$$

is the so called Householder matrix, and  $\delta_{\alpha\beta}$  is the identity matrix. If  ${}^jV_{\alpha i}$  is premultiplied by  ${}^j B_{\alpha\beta}^*$ , a new matrix  ${}^j B_{\alpha\beta}^* {}^jV_{\beta i}$  is obtained:

$${}^1 B_{\alpha\beta}^* {}^1V_{\alpha i} = \begin{bmatrix} \|v_{\alpha 1}\| & \cdots & \cdots & \cdots \\ 0 & & & \\ \vdots & & {}^2V_{\beta j} & \\ 0 & & & \end{bmatrix} \quad (19)$$

Matrix 19 is upper triangular in the first  $j$  columns; and dense everywhere else. A new submatrix  ${}^{j+1}V_{\beta j}$  is therefore defined after erasing the first column and row. This process can be repeated until the initial matrix becomes upper triangular.

Once the algorithm is computed for every column on  ${}^1V_{\alpha i}$ , a set of  $q$  gradually smaller matrices  ${}^1B_{\alpha i}^* \in \mathbb{R}^{p \times p}$ ,  ${}^2B_{\alpha i}^* \in \mathbb{R}^{p-1 \times p-1}$  ...  ${}^jB_{\alpha i}^* \in \mathbb{R}^{p-(j-1) \times p-(j-1)}$  ...  ${}^qB_{\alpha i}^* \in \mathbb{R}^{1 \times 1}$  are obtained. To properly compute the  $j$ -ith iteration of eq. (19), matrices  ${}^jB_{\alpha i}$  are completed with the identity:

$${}^jB_{\alpha i} = \begin{bmatrix} I_{ij} & 0 \\ 0 & {}^jB_{\alpha i}^* \end{bmatrix} \quad (20)$$

where  $I_{ij} \in \mathbb{R}^{j-1 \times j-1}$ . Finally, through eqs. (14) and (15) the matrices  $U_{ij}$  and

$Q_{\alpha i}$  are obtained. The process is described in algorithm 2.

```

1  ${}^1V_{\alpha i} = V_{\alpha i}$ 
  for  $j=1\dots q$  do
2   choose  $v_\alpha = {}^jV_{\alpha i}$  with  $\alpha = j\dots p$  and  $i = j$ 
3    $n_\alpha = v_\alpha - \|v_\alpha\|e_\alpha$ 
4    $u_\alpha = n_\alpha/\|n_\alpha\|$ 
5    $B_{\alpha\beta}^* = I - 2 u_\alpha u_\beta$ 
6    ${}^{j+1}V_{\alpha i} = B_{\alpha\beta} V_{\beta i}$ 
  end
7  $Q_{\alpha\epsilon} = {}^1B_{\alpha\beta} {}^2B_{\beta\gamma}\dots {}^qB_{\gamma\epsilon}$ 
8  $U_{\alpha i} = {}^qB_{\alpha\beta} {}^{q-1}B_{\beta\gamma}\dots {}^1B_{\gamma\epsilon} V_{\epsilon i}$ 

```

**Algorithm 2:** overview of the QR decomposition algorithm.

### 2.5. Paralell compact IQN

The distributed memory parallelisation of Alya is based on a domain decomposition [25], a mesh partition is carried out [26], and each partition is assigned to a MPI-process. The mesh partitioner divides the mesh minimising the area between subdomains but without any requirements on the contact surface  $\Gamma_c$  (see fig. 1). Therefore, the nodes in  $\Gamma_c$  will be distributed among the MPI tasks and so the increment matrix  $V_{\alpha i}$ .

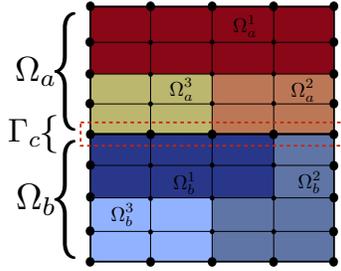


Figure 1: Physical subdomains  $\Omega_a$  and  $\Omega_b$  in contact by the wet surface  $\Gamma_c$ . Each physical subdomain is subdivided in three computational subdomains (partitions). The wet surface, and therefore the vectors in matrix  $V_{\alpha i}$  can be distributed along several partitions.

In order to process  $V_{\alpha i}$  in parallel, we look for the subdomain with the largest number of contact nodes, the now called “leader” partition. After,  $V_{\alpha i}$  and  $W_{\alpha i}$  are renumbered so the first rows correspond to the leader partition, so the backsubstitution is only executed there.

To improve computing and memory cost, we propose some modifications for the base algorithm in section 2.4. As  ${}^j B_{\alpha\beta}$  is obtained by eq. (18), the product  ${}^j B_{\alpha\beta} {}^j V_{\beta i}$  can be expanded as:

$${}^j B_{\alpha\beta} {}^j V_{\beta i} = (\delta_{\alpha\beta} - 2u_{\alpha}u_{\beta}) {}^j V_{\beta i} = {}^j V_{\alpha i} - 2u_{\alpha}u_{\beta} {}^j V_{\beta i}. \quad (21)$$

So, instead of computing and storing  ${}^j B_{\alpha\beta} \in \mathbb{R}^{p \times p}$  for each iteration  $j$ , we store the vectors  $u_{\alpha} \in \mathbb{R}^p$  for the  $q$  iterations. Expanding eq. (21):

$$(\delta_{\alpha\beta} - 2u_{\alpha}u_{\beta}) {}^j V_{\beta i} = \delta_{\alpha\beta} {}^j V_{\beta i} - 2u_{\alpha}u_{\beta} {}^j V_{\beta i} = {}^j V_{\alpha i} - 2u_{\alpha}(u_{\beta} {}^j V_{\beta i}). \quad (22)$$

So the parallel matrix-vector product  $u_{\alpha} {}^j V_{\alpha i}$  can be computed first, then compute  $u_{\alpha}(u_{\beta} {}^j V_{\beta i})$  and finally subtract  $I_{\alpha\beta} {}^j V_{\beta i} - 2u_{\alpha}u_{\beta} {}^j V_{\beta i}$ . To compute eq. (15) we proceed similarly, but starting with  $(\delta_{\alpha\beta} - 2{}^1 u_{\alpha} {}^1 u_{\beta})V_{\beta i}$  followed by the premultiplication of the matrices  $(\delta_{\alpha\beta} - 2u_{\alpha}u_{\beta})$  with the same technique as described here.

Similarly, we can avoid the construction of the dense matrix  $Q_{\alpha\beta} \in \mathbb{R}^{p \times p}$ , used in the backsubstitution (see eq. (10)). Vector  $-Q_{\alpha i}r_{\alpha}$  can be computed with a strategy similar to eq. (21). The difference is that  $Q_{\alpha\epsilon} = {}^1 B_{\alpha\beta} \cdots {}^q B_{\gamma\epsilon}$ , so after computing  ${}^q B_{\alpha i}r_{\alpha}$  as:

$${}^q B_{\alpha\beta}r_{\alpha} = (\delta_{\alpha\beta} - 2{}^q u_{\alpha} {}^q u_{\beta})r_{\beta} = r_{\beta} - 2{}^q u_{\alpha} {}^q u_{\beta}r_{\beta}, \quad (23)$$

the rest of the matrices  $B_{\alpha\beta} = \delta_{\alpha\beta} - 2u_{\alpha}u_{\beta}$  are premultiplied. The first multiplication, (eq. (23)) can be expanded as:

$$(\delta_{\alpha\beta} - 2{}^q u_{\alpha} {}^q u_{\beta})r_{\beta} = \delta_{\alpha\beta}r_{\beta} - 2{}^q u_{\alpha} {}^q u_{\beta}r_{\beta} = r_{\alpha} - 2{}^q u_{\alpha} ({}^q u_{\beta}r_{\beta}) \quad (24)$$

The product  ${}^q u_{\beta}r_{\beta}$  is firstly computed and then  $r_{\alpha} - 2{}^q u_{\alpha} ({}^q u_{\beta}r_{\beta})$ . The resulting vector is multiplied by  $(\delta_{\alpha\beta} - 2({}^{q-1})u_{\alpha} ({}^{q-1})u_{\beta})$  and followed by every

matrix  $(\delta_{\alpha\beta} - 2^j u_\alpha^j u_\beta^j)$  up to  $(\delta_{\alpha\beta} - 2^1 u_\alpha^1 u_\beta^1)$ . In this way, matrices  ${}^q B_{\alpha i}$  are never completely computed.

The resulting algorithm has as input the  $V_{\alpha i}$  matrix and the residuals vectors to operate in the backsubstitution (see eq. (10)), and the output will be the coefficient vector  $\alpha_i$ . As the boundaries between the IQN and the QR algorithms can't be identified anymore, we refer to the developed algorithm as Compact IQN (CIQN). Our main motivation is that a complete QR decomposition would be prohibitive in large cases as a dense  $Q$  orthogonal matrix would be extremely expensive to compute and store. The proposed algorithm is a collection of matrix-vector and vector-vector products restricted to the contact. An efficient parallelisation requires a proper point-to-point MPI communication on the modified IQN and QR algorithms. The whole sequence of steps is described

in algorithm 3.

```

1 Chose leader partition
  while not the last time step do                                     ▷ time loop
2    $x_\alpha^0 = x_\alpha^{ini}$ 
3    $\tilde{x}_\alpha^0 = H(x_\alpha^0)$ 
4    $r_\alpha^0 = \tilde{x}_\alpha^0 - x_\alpha^0$ 
5   update  $x_\alpha^1 = \tilde{x}_\alpha^0 + \omega_0 r_\alpha^0$ 
   while problem not converged do
6      $\tilde{x}_\alpha = H(x_\alpha)$ 
7      $r_\alpha = \tilde{x}_\alpha - x_\alpha$ 
8     build  $V_{\alpha i}$  and  $W_{\alpha i}$  as eqs. (12) and (13)
9      ${}^1V_{\alpha i} = V_{\alpha i}$ 
   for  $j=1\dots q$  do                                               ▷ QR decomposition loop
10    if  $j > 1$  then  ${}^{j+1}V_{\alpha i} = {}^1V_{\beta i} - 2u_\alpha u_\beta {}^1V_{\beta i}$  as eq. (22)
11    if I am the leader then
12    |  $v_\alpha = {}^jV_{\alpha i}$  with  $\alpha = j\dots p$  and  $i = j$ 
13    else
14    |  $v_\alpha = {}^jV_{\alpha i}$  with  $\alpha = 1\dots p$  and  $i = j$ 
15    end
16    parallel compute  $\|v_\alpha\|$ 
17    if I am leader then  $n_\alpha = v_\alpha - \|v_\alpha\| {}^j e_\alpha$ 
18    parallel compute  $\|n_\alpha\|$ 
19     $u_\alpha = n_\alpha / \|n_\alpha\|$ 
20    end
21     $U_{ij} = (\delta_{i\gamma} - 2 {}^q u_i {}^q u_\gamma) \dots ({}^1V_{\alpha j} - 2 {}^1u_\alpha {}^1u_\beta {}^1V_{\beta j})$  (as eq. (22))
22    if  $|U_{jj}| < \epsilon \|U\|_2$  then remove column  $j$  and restart QR loop
23     $-Q_{\alpha i} r_\alpha = -(\delta_{i\gamma} - 2 {}^1u_i {}^1u_\gamma) \dots (r_\beta - 2 {}^q u_\beta {}^q u_\alpha r_\alpha)$  (as eq. (24))
24    if I am leader then backsubstitute  $U_{ij} \lambda_j = -Q_{\alpha i} r_\alpha$ 
25     $x_\alpha^{J+1} = \tilde{x}_\alpha + W_{\alpha i} \lambda_i$                                ▷ variable update
  end
end

```

**Algorithm 3:** Compact Interface quasi-Newton algorithm.

## 2.6. Physics of solved cases

The algorithm in this work has been developed generically for any interface problem. Although that, the main interest of the authors is Fluid-Structure Interaction (FSI), we briefly describe the governing equations. The Newtonian fluid is modelled with incompressible Navier-Stokes equations using an Arbitrary Lagrangian-Eulerian (ALE) formulation:

$$\rho^f \frac{\partial u_i}{\partial t} + \rho^f (u_j - u_j^m) \frac{\partial u_i}{\partial x_j} + \frac{\partial}{\partial x_j} \left[ +p\delta_{ij} - \mu \left( \frac{\partial u_i}{\partial x_j} + \frac{\partial u_j}{\partial x_i} \right) \right] = +\rho^f f_i \quad (25)$$

$$\frac{\partial u_i}{\partial x_i} = 0, \quad (26)$$

where  $\mu$  is the viscosity of the fluid,  $\rho^f$  the density,  $u_i$  the velocity,  $p$  is the mechanical pressure,  $f_i$  the force term and  $u_j^m$  is the fluid domain velocity. The numerical model is based on the Finite Element Method, using the Variational Multiscale[27]. For the Arbitrary Lagrangian-Eulerian (ALE) formulation, the technique used is proposed in [28]. Mesh movement is solved through a Laplacian equation

$$\frac{\partial}{\partial x_j} \left( [1 + \alpha^e] \frac{\partial b_i}{\partial x_j} \right) = 0, \quad (27)$$

where  $b_i$  are the components of the displacement in each point for the domain. The factor  $\alpha$  is a diffusive term that, once discretised, controls the mesh distortion. ALE boundary conditions at the contact surface is set through the nodal displacement from solid mechanics problem.

Solid mechanics is solved following a transient scheme and using a total Lagrangian formulation in finite strains [13]. The displacement  $d_i$  form of the linear momentum balance can be modelled as:

$$\rho^s \frac{\partial^2 d_i}{\partial t^2} = \frac{\partial P_{iJ}}{\partial X_J} + \rho^s B_i, \quad (28)$$

where  $\rho^s$  is the initial density of the body,  $B_i$  represents the body forces and  $P_{iJ}$  is the nominal stress tensor. Solid mechanics boundary conditions at the contact surface is set through the nodal forces from the fluid mechanics problem.

Let us label ‘‘CFD’’ and ‘‘CSM’’ the fluid and solid sides of a coupled FSI problem. At the contact surface, displacements and normal stresses must be

continuous:

$${}^{\text{CFD}}d_i^{\Gamma_c} = {}^{\text{CSM}}d_i^{\Gamma_c} \quad (29)$$

$$n_i {}^{\text{CFD}}\sigma_{ij}^{\Gamma_c} = n_i {}^{\text{CSM}}\sigma_{ij}^{\Gamma_c}, \quad (30)$$

where  ${}^{\text{CFD}}d_i^{\Gamma_c}$  and  ${}^{\text{CSM}}d_i^{\Gamma_c}$  are the deformation in the contact boundary for the fluid and for the solid respectively; and  $n_i {}^{\text{CSM}}\sigma_{ij}^{\Gamma_c}$  and  $n_i {}^{\text{CFD}}\sigma_{ij}^{\Gamma_c}$  are the normal stresses in the contact boundary.

A typical behaviour of the developed algorithm is shown in Appendix C.1 for an FSI case. Solver iteration are decomposed in the different parts: Momentum, continuity and ALE for the fluid domain, and displacement for the solid domain.

### 3. Results and discussion

In this section we present three cases. Problems in sections 3.2 and 3.3 were chosen to show the difference of the behaviour of the algorithm with different dynamics on the physics, while problem in section 3.4 is an scalability test. Every problem in this section can be executed relaxing the force or the displacement. The best (less average iterations) scheme for each case is shown in this section, the rest in Appendix C. For each case, a sensitivity analysis is executed for the number of past saved time steps (histories), iterations on each time step (ranking) and  $\epsilon$ . Also, as a reference, the number of iterations is compared against the popular Aitken algorithm. For this comparison algorithm, results are shown as (*e.g.*) 17.76 (sd=2.91) where the first figure indicates the mean and the second figure the standard deviation (sd). All cases are executed in Marenstrum IV supercomputer.

#### 3.1. Algorithm validation

Validating the algorithm is a mandatory step to trust the results in this section. The numerical method is validated with the benchmark **FSI3** proposed in section 4.3 of [29]. The experimental set-up involves a flexible rod oscillating in a fluid flow. The dimensions of the fluid domain are  $41.0 \times 250.0[cm]$ ,

the dimensions of the rod  $2.0 \times 35.0 [cm]$  and the radius of the anchoring structure being  $r = 5.0 [cm]$ . The densities for the fluid and the solid are  $\rho_f = \rho_s = 1 [g/cm^3]$ . The dynamic viscosity for the fluid is  $\mu_f = 10.0 [Poise]$  and the Young modulus and Poisson's ratio for the isotropic solid are  $E = 5.6E7 [Barye]$  and  $\nu = 0.4 [-]$  respectively. The fluid and solid meshes are composed by  $35k$  and  $16k$  linear triangles respectively and the problem solved with a time step of  $1E-3 [s]$ . Oscillation frequency and amplitude at the tip of the rod are measured to compare against numerical results. Results for  $t=3 [s]$  are shown in fig. 2.

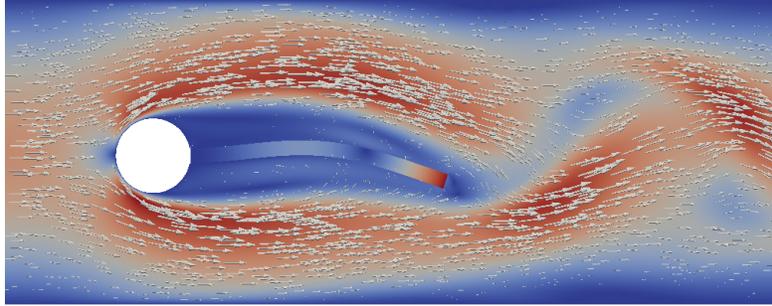


Figure 2: Method Validation. Portion of the domain proposed by [29] in time  $t = 3 [s]$ . Deformation is represented on the bar and velocity field in the fluid domain.

For our code, the obtained amplitude and frequency on the quasi-periodic period are  $A_x = -2.60 \times 10^{-3} \pm 2.40 \times 10^{-3} [f = 10.8]$  and  $A_y = 2.3 \times 10^{-3} \pm 33.7 \times 10^{-3} [f = 5.4]$  in concordance with the  $A_x = -2.69 \times 10^{-3} \pm 2.53 \times 10^{-3} [f = 10.9]$  and  $A_y = 1.48 \times 10^{-3} \pm 34.38 \times 10^{-3} [f = 5.3]$  obtained in the cited experiment<sup>1</sup>. Although there is already a good agreement, results can be further improved by refining the meshes, as proven in section 4.1 of [30]. With this, we prove the algorithm is correctly implemented and reproduces the physics of the FSI problem.

<sup>1</sup>The results are presented as in the original experiment: mean  $\pm$ amplitude[*freq*].

### 3.2. Wave propagation in elastic tube

The domain, schematised in fig. 3, is an elastic tube, filled with fluid. The densities are  $\rho_f = \rho_s = 1[g/cm^3]$  for the fluid and the solid. Fluid viscosity is  $\mu = 0.03[Poise]$ . The Young modulus and Poisson's ratio for the solid are  $E = 3E7[Baryes]$  and  $\nu = 0.3[-]$ . Inflow velocity is  $30[cm/s]$ . Outflow pressure is  $p = 0[Baryes]$ . In the contact surface, continuity of displacement and normal stresses are imposed. Linear tetrahedra are used for the spatial discretisation resulting in 48k elements (10k nodes) and 30k elements (6k nodes) for the fluid and the solid respectively, with 2.7k interface nodes ( $\sim 25\%$  of the total). Time step is fixed at  $\Delta t = 4E-4[s]$ . For each case a sensitivity analysis is done with a range of previous time steps, iterations and  $\epsilon$ . Each case run in 24 cores in Marenostrum IV. As a reference for the reader, with the optimal configurations the CIQN algorithm case took 38 minutes and the Aitken case took 67 minutes 13 seconds.

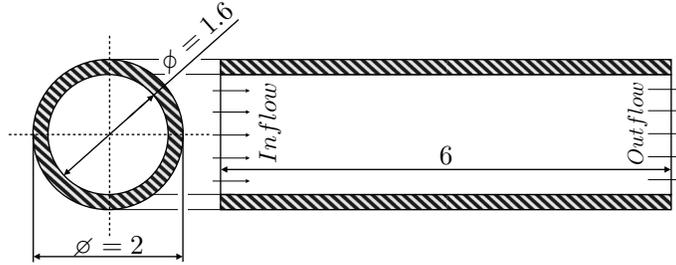


Figure 3: Scheme with dimensions for the wave propagation in elastic tube experiment.

Results and statistical tendencies relaxing displacement are shown in table 1 and fig. 4. Similar information, but relaxing force is shown in Appendix C.2. For a qualitative comparison, with the Aitken algorithm, the solver requires 17.76 (sd=2.91) and 21.81 (sd=3.64) iterations in average when relaxed on force and displacement respectively.

histories	ranking	$\epsilon=0$	$\epsilon=1E-9$	$\epsilon=1E-7$	$\epsilon=1E-5$	$\epsilon=1E-3$	$\epsilon=0.1$
0	5	13.36	13.36	13.36	13.24	13.08	17.66
	10	12.32	12.32	12.30	12.82	15.16	28.98
1	5	11.25	11.02	12.86	13.24	13.08	17.66
	10	11.72	11.88	12.41	12.92	15.16	28.97
2	5	10.36	10.55	12.82	13.24	13.08	17.66
	10	11.43	11.57	11.77	12.92	15.16	28.97
5	5	9.67	10.62	12.87	13.24	13.08	17.66
	10	11.59	11.82	12.31	12.89	15.16	28.97
10	5	9.64	10.49	12.84	13.24	13.08	17.66
	10	11.64	11.68	12.31	12.92	15.16	28.97

Table 1: Results for the wave propagation in an elastic tube experiment when displacement is relaxed. Iterations for the scheme depending on the number of previous time steps used (histories), iterations in each time step (ranking) and filter ( $\epsilon$ ) when relaxing displacement. The average number of iterations is 14.2925 (sd=4.72).

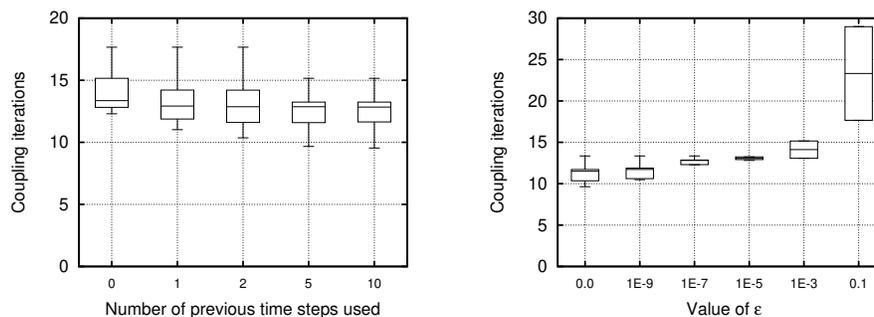


Figure 4: Candle plots for the wave propagation in an elastic tube experiment when displacement is relaxed. Less coupling iterations are required with a larger number of previous time steps and a smaller  $\epsilon$ .

Although simple, the problem in this section has been similarly reproduced in other FSI articles [16, 31, 21, 19]. Here we show that, as similarly concluded in [19] adding information from previous time steps improves the rate of con-

vergence of the algorithm (left plot on fig. 4). On the contrary, increasing the ranking does not necessarily have a positive effect. Filtering has the effect of reducing the standard deviation on the number of iteration in each time step (right plot on fig. 4), but it is arguable if it compensates the added computational cost of restarting the QR decomposition. A similar behaviour can be seen if forced is relaxed (see Appendix C.2).

### 3.3. Oscillating rod and flexible wall in a fluid domain

The domain, schematised in fig. 5, is composed by a centered oscillating flexible rod and a fixed flexible wall, being the rest the fluid domain. The densities for the fluid and the solid are  $\rho_f = \rho_s = 1[g/cm^3]$ . Fluid viscosity is  $\mu = 0.04[Poise]$ . The Young modulus and poisson ratio for the solid are  $E = 2E7[Baryes]$  and  $\nu = 0.3[-]$ . The tip of the oscillating rod has an imposed displacement of  $d_x = \sin(2\pi t)$ . The Inflows velocity is  $0.1[cm/s]$ . Outflows pressures are  $p = 0[Baryes]$ . Continuity of displacement and normal stresses are imposed in both contact surfaces. Linear triangles are used for the spatial discretisation resulting in 7.4k elements (4k nodes) and 752 elements (474 nodes) for the fluid and the solid respectively, with 300 interface nodes ( $\sim 10\%$  and  $\sim 45\%$  of the total for the fluid and the solid respectively). Time step is fixed at  $\Delta t = 0.1[s]$ . Each case run in 16 cores in Marenstrum IV. As a reference for the reader, with the optimal configurations, the CIQN algorithm case took 57 seconds and the Aitken case took 7 minutes and 49 seconds.

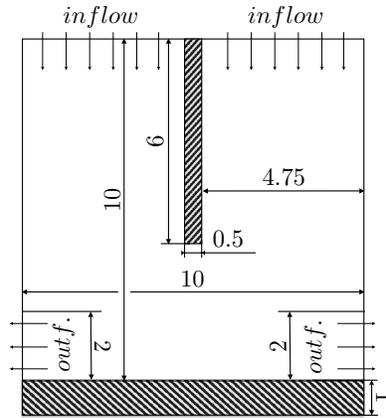


Figure 5: Scheme with dimensions for the oscillating rod with flexible wall experiment.

Results and statistical tendencies relaxing displacement are shown in table 2 and fig. 6. Similar information, but relaxing force is shown in Appendix C.3. For a qualitative comparison, with the Aitken algorithm, the solver requires 55.46 (sd=41.43) and 69.98 (sd=55.94) iterations in average when relaxed on force and displacement respectively.

histories	ranking	$\epsilon=0$	$\epsilon=1E-9$	$\epsilon=1E-7$	$\epsilon=1E-5$	$\epsilon=1E-3$	$\epsilon=0.1$
0	5	12.02	12.02	12.02	11.96	F	13.44
	10	13.94	13.94	13.88	13.82	12.68	17.98
1	5	14.56	14.52	13.14	12.36	11.66	13.44
	10	11.8	11.74	F	13.76	12.52	17.98
2	5	12.92	12.4	12.76	12.32	11.66	13.44
	10	10.62	10.92	12.76	13.76	12.52	17.98
5	5	F	11.9	13.42	12.32	11.66	13.44
	10	10.42	10.56	13.94	13.76	12.52	F
10	5	14.30	11.96	13.42	12.32	11.66	13.44
	10	12.22	12.34	13.94	13.76	12.52	17.98

Table 2: Results for the oscillating rod experiment when displacement is relaxed. Iterations for the scheme depending on the number of previous time steps used (histories), iterations in each time step (ranking) and filter ( $\epsilon$ ) when relaxing displacement, F meaning a diverged simulation. The average number of iterations is 14.99 (sd=8.33).

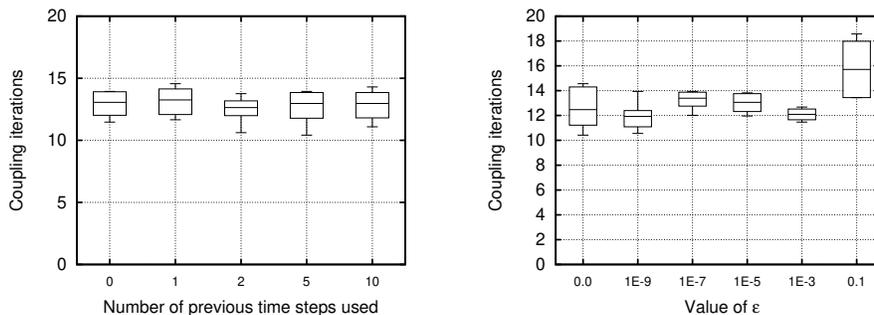


Figure 6: Candle plots for the oscillating rod experiment when displacement is relaxed. The number of coupling iterations do not seem to improve with filtering or re-usage of previous information.

The problem presented in this section, even though being computationally cheaper, is more physically challenging. Compared to problem in section 3.2, there are two surfaces to couple and the deformations are considerably larger.

This is reflected in the computational cost, requiring 55.46 (sd=41.43) Aitken iterations <sup>2</sup>. In this case, the behaviour of the IQN algorithm is completely different from the presented in section 3.2. Now we can see scattered diverging simulations in table 2, increasing the number of histories do not bring any advantage (left side of fig. 6) and using filtering doesn't bring any drastic improvement (right side of fig. 6). On the contrary, and oppositely to the case presented in section 3.2, increasing the rank has a beneficial effect on the convergence properties of the algorithm when little or no filtering is included. The strongest hypothesis for this behaviour is that including histories over-predicts the final position of the interface, hampering the convergence of the algorithm.

### 3.4. Scalability

The domain, schematised in fig. 7, is a filled flexible tube lying over a flexible surface which is in contact with a big volume of another fluid. The densities for the fluid and the solid are  $\rho_f = \rho_s = 1[g/cm^3]$ . Fluid viscosity is  $\mu = 0.03[Poise]$ . The Young modulus and Poisson ratio for the solid are  $E = 1.5E4[Baryes]$  and  $\nu = 0.3[-]$ . The Inflow velocities are  $\sin(2\pi t)$  and  $1[cm/s]$  for the lower domain respectively. Outflows pressures are  $p = 0[Baryes]$ . Continuity of displacement and normal stresses are imposed in both contact surfaces. Linear tetrahedra are used for the spatial discretisation resulting in 60M elements (10.4M nodes) and 40M elements (7.1M nodes) for the fluid and the solid respectively, with 4M interface nodes ( $\sim 38\%$  and  $\sim 56\%$  of the total nodes for the fluid and the solid respectively). Time step is fixed at  $\Delta t = 0.1[s]$ .

---

<sup>2</sup>compared to the 17.76 (sd=2.91) Aitken iterations required by the elastic tube experiment.

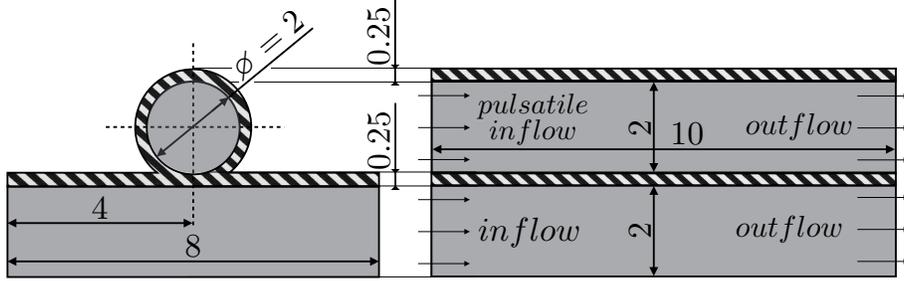


Figure 7: Scheme with dimensions for the scalability case.

As the main goal is to prove good scalability, the results for the sensibility analysis are shown in Appendix C.4. Table 3 show results for the solver running independently (uncoupled) to ease its comparison with the coupled scalability. To obtain an optimum efficiency  $E$ , several cases are run sweeping the core allocation for each physical problem. Figure 8 shows speed-up and efficiency for four fixed values in the fluid solver core count  $p_f$ . In each case the core count for the solid mechanic solver  $p_s$  is ranged between 64 and 2048, with increments in power of two. This processes is performed for a core count of 256, 512, 1024, and 2048 in the fluid solver. The results is a set of curves with a peak efficiency  $E_{p_f}$  given by the optimal balance of cores for each case.

Core count	Fluid mechanics		Solid mechanics	
	speed up	efficiency	speed up	efficiency
128	128.0	1.00	128.0	1.00
256	256.0	0.99	256.0	0.99
512	511.6	0.99	508.1	0.99
1024	1011.0	0.98	960.5	0.93
2048	1880.3	0.91	1793.0	0.87

Table 3: Parallel performance analysis. Speed-up and efficiency for both solvers.

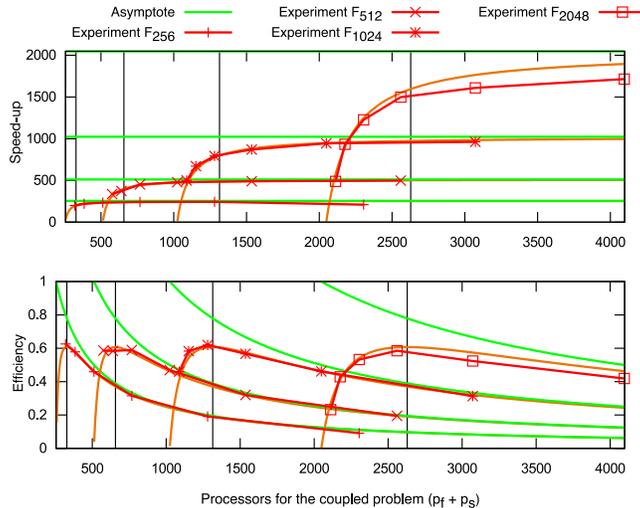


Figure 8: Parallel performance analysis. Speed-up and Efficiency for a core allocation of  $p = p_f + p_s$ , where  $p_f = \{256, 512, 1024, 2048\}$  in the core count for the fluid, and  $p_s = \{64, 128, 256, 512, 1024, 2048\}$  for the solid. The orange line is the fitted curve for each case. An optimal allocation  $p_{opt}$  which allows to achieve the maximum efficiency  $E_{p_f}$  of the coupled system can be found for each curve  $S_{p_f}$ .

This example demonstrates the necessity of the parallel version of the algorithm, as it would be impossible to fit 11M interface nodes in a single node of a shared-memory high-performance computing infrastructure. Furthermore, we show a scalability up to 4800 cores. Although the scalability of the uncoupled problem do not drop under 87% in the uncoupled case, when coupled, the maximum scalability achieved is 60%. This is due to the staggered scheme used that improves stability, one set of cores is idle while the rest is processing. Although it's out of the scope of this article, this issue can be tackled by using core overloading on the MPI scheme.

For another large scale use of the algorithm, please refer to [32], where a human heart is solved with the IQN algorithm.

#### 4. Conclusions & future work

In this paper, we introduced the compact interface quasi-Newton (CIQN) coupling scheme, optimised for distributed memory architecture. The developed algorithm includes reusing information from the previous time steps (histories) and filtering, but in an efficient scheme that avoids constructing dense matrices and reduces the number of operations. This leads to an algorithm that requires less coupling iterations and computing time per time step.

In previous works [18, 19] it has been stated that using information from previous time steps together with filtering improve convergence. In this work we show this do not necessarily happen and can only be stated for certain type of dynamic behaviour, while a correct parametrisation requires a fine tuning for each case.

In this work we prove that a parallel coupling algorithm is mandatory for massively parallel cases. We also show that reusing histories does not necessarily improves the convergence rate of the algorithm, but is dependant on the dynamics of the problem. As it has been said, there is no a silver bullet algorithm to tackle all the FSI cases. The chosen algorithm and parameters must fit the features of the problem to solve, taking into account the dynamics and the possible numerical instabilities that may arise.

Although the developed algorithm has been proved robust and efficient, there is room for optimise the execution with core-overloading. Also, other filtering algorithms (*e.g.* eigenvalue-based) can be tested to compare performance and computational cost. Finally, the behaviour of the presented algorithm has to be tested in other interface problems like solid-solid contact or heat transmission. These topics will be developed in a future work.

#### Acknowledgements

This work has been funded by CompBioMed project a grant by the European Commission H2020 (agreement nr: 823712), EUBrazilCC a project under the

Programme FP7-ICT (agreement nr: 614048) and a FPI-SO grant (agreement nr: SVP-2014-068491).

## Appendix A. Glossary

Glossary of acronyms used at the manuscript.

- **ALE:** Arbitrary Lagrangian-Eulerian.
- **BSC:** Barcelona Supercomputing Center.
- **CFD:** Computed Fluid Dynamics.
- **CIQN:** Compact interface quasi-Newton.
- **CSM:** Computed Solid Mechanics.
- *E*: Young modulus.
- **FSI:** Fluid-structure interaction.
- **HPC:** High Performance Computing.
- **IQN:** Interface quasi-Newton.
- **MPI:** Message Passing Interface.
- **sd:** Standard deviation.
- $\nu$ : Poisson's ratio.
- $\rho$ : Density.
- $\mu$ : Dynamic viscosity.

## Appendix B. Index notation convention

To ease implementation, the Einstein convention on repeated indices will be followed, allowing to describe the mathematics, the physics and the computational implementation aspects depending on the context. For the continuum problem, the indices label space dimensions. On the discretised problem, the lowercase greek alphabet  $\alpha = 1, \dots, p$  labels the number of degrees of freedom  $p$ , i.e. the matrix rows. The lowercase latin alphabet labels the matrix columns,  $i = 1, \dots, q - 1$  where  $q$  is the last stored iteration. Additionally, a capital latin subindex labels the FSI solver iteration  $I = 1, \dots, q - 1$ , where  $q$  is the last stored iteration. A final rule is how those indices operate: only those of the same kind are contracted. For instance,  $Q_{\alpha i}^{I-1}$  is the  $Q$  matrix for iteration  $I - 1$  with rows labelled  $\alpha$  and columns  $i$ . When this matrix is multiplied by a certain vector  $B_i$ , it results in a given vector  $A_{\alpha}^{I-1}$ :

$$A_{\alpha}^{I-1} = Q_{\alpha i}^{I-1} B_i = \sum_{i=1}^{q-1} Q_{\alpha i}^{I-1} B_i,$$

where latin indices  $i$  are contracted.

## Appendix C. More results

### *Appendix C.1. Typical behaviour of the coupling residue and solver iterations*

In Figure C.9 we show the typical behaviour of the coupling residue and solver iterations for a single time step. Particularly, it corresponds to the 20th time step of the case presented in section 3.3 using 5 iterations per time step (ranking=5) and 5 previous time iterations (history=5) and an  $\epsilon=1\text{E-}9$ .

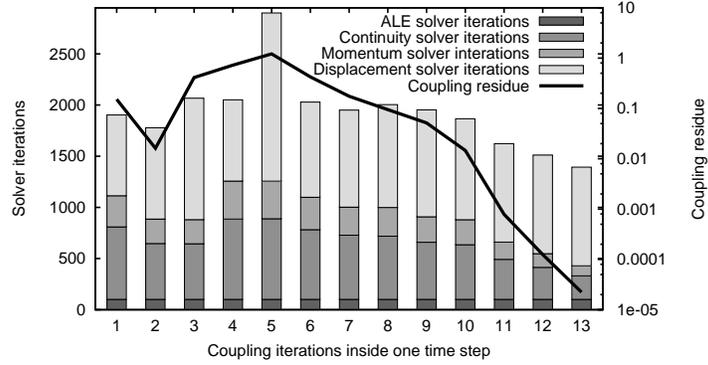


Figure C.9: Iterations of the different solvers and residue of the coupling for a typical time step in the CIQN algorithm.

*Appendix C.2. For the wave propagation in elastic tube*

Table C.4 and fig. C.10 show the results when force is relaxed.

histories	ranking	$\epsilon=0$	$\epsilon=1E-9$	$\epsilon=1E-7$	$\epsilon=1E-5$	$\epsilon=1E-3$	$\epsilon=0.1$
0	5	11.89	11.84	11.79	11.80	12.22	18.12
	10	15.24	15.29	15.55	15.13	17.19	33.90
1	5	11.42	11.54	11.83	11.76	12.22	18.12
	10	14.46	15.07	15.60	15.18	17.11	33.9
2	5	11.33	11.87	11.76	11.12	12.22	18.12
	10	13.99	15.04	15.78	15.20	17.11	33.90
5	5	11.87	10.68	11.34	11.76	12.22	18.12
	10	13.80	14.89	15.78	15.20	17.11	26.74
10	5	10.59	11.45	11.87	11.76	12.22	18.12
	10	13.77	14.73	15.78	15.20	17.11	33.90

Table C.4: Results for the wave propagation in an elastic tube experiment when relaxing force. Iterations for the scheme depending on the number of previous time steps used (histories), iterations in each time step (ranking) and filter ( $\epsilon$ ) when relaxing displacement. caption relaxing force. The average number of iterations is 15.298125(sd=5.58)

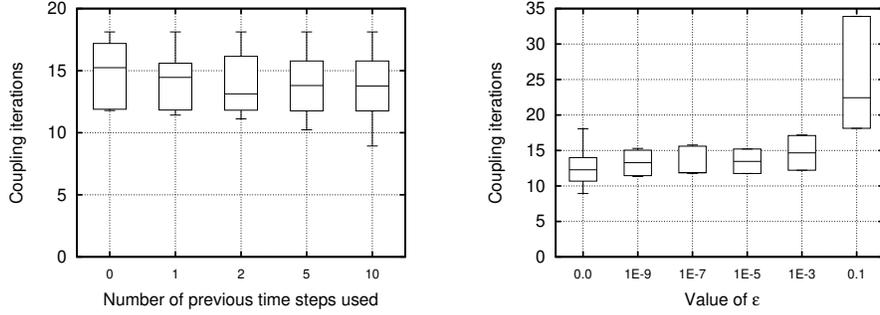


Figure C.10: Candle plots for the wave propagation experiment when force is relaxed.

*Appendix C.3. For the oscillating bar and flexible wall in a fluid domain*

Table C.5 and fig. C.11 show the results when force is relaxed.

histories	ranking	$\epsilon=0$	$\epsilon=1E-9$	$\epsilon=1E-7$	$\epsilon=1E-5$	$\epsilon=1E-3$	$\epsilon=0.1$
0	5	13.78	13.78	13.78	13.84	13.78	15.78
	10	19.48	19.48	19.30	18.72	16.84	24.42
1	5	13.02	13.22	13.38	13.72	13.78	15.78
	10	16.42	16.42	17.14	18.24	17.48	24.44
2	5	13.46	13.46	13.82	13.72	13.78	15.78
	10	F	F	16.96	18.26	17.48	24.44
5	5	12.94	13.10	13.24	13.72	13.78	15.78
	10	15.60	15.14	15.62	18.26	17.48	24.44
10	5	12.20	11.66	13.42	13.72	13.78	15.78
	10	15.28	14.91	15.62	18.26	17.48	24.44

Table C.5: Results for oscillating rod experiment when force is relaxed. Iterations for the scheme depending on the number of previous time steps used (histories), iterations in each time step (ranking) and filter ( $\epsilon$ ) when relaxing displacement, F meaning a divergent case. The average number of iterations is 16.04 (sd=3.19).

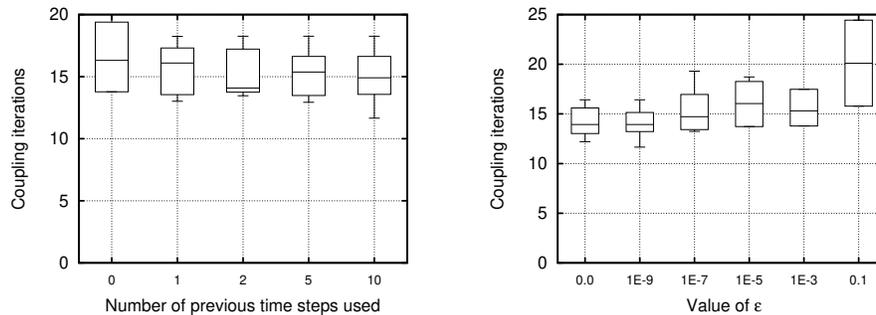


Figure C.11: Candle plots for the oscillating rod experiment when force is relaxed.

#### Appendix C.4. For the scalability

More results for the scalability case. As the objective of this experiment is to prove the HPC performance of the algorithm all the numerical sensitivity analysis is shown here. Figure 7 shows a scheme of the geometry used. Table C.6 and fig. C.12 show results when displacement is relaxed and table C.7 and fig. C.13 when force is relaxed. For a quantitative comparison, with the Aitken algorithm the scheme requires 19.45 (sd=3.10) and 19.58 (sd=2.16) when relaxing on force and displacement respectively. Each case run in 768 cores in Marenstrum IV. As a reference for the reader, in the optimal configurations, the CIQN algorithm case took 1 hour, 27 minutes and 57 seconds and the Aitken case took 2 hours, 27 minutes and 30 seconds.

histories	ranking	$\epsilon=0$	$\epsilon=1E-9$	$\epsilon=1E-7$	$\epsilon=1E-5$	$\epsilon=1E-3$	$\epsilon=0.1$
0	5	14.85	14.85	14.85	14.83	13.80	19.62
	10	15.39	15.39	15.41	15.57	14.69	32.56
1	5	12.03	12.03	12.15	13.31	13.91	19.63
	10	14.31	14.31	14.46	14.98	14.79	32.97
2	5	11.38	11.38	11.87	12.98	13.91	19.63
	10	13.88	13.56	13.94	15.38	14.79	32.97
5	5	11.81	13.23	11.51	12.97	13.91	19.63
	10	12.75	13.35	13.73	15.38	14.79	32.97
10	5	11.85	12.69	12.11	12.97	13.91	19.63
	10	12.70	12.87	13.99	15.38	14.79	32.97

Table C.6: Results for scalability case when relaxing displacement. Iterations for the scheme depending on the number of previous time steps used (histories), iterations in each time step (ranking) and filter ( $\epsilon$ ) when relaxing displacement. The average number of iterations is 15.80 (sd=5.52).

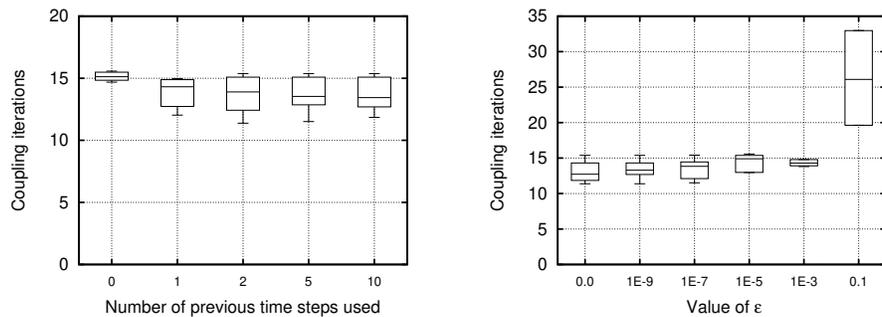


Figure C.12: Candle plots for the scalability when displacement is relaxed.

histories	ranking	$\epsilon=0$	$\epsilon=1E-9$	$\epsilon=1E-7$	$\epsilon=1E-5$	$\epsilon=1E-3$	$\epsilon=0.1$
0	5	13.40	13.40	13.40	13.36	13.26	20.33
	10	20.50	20.42	19.95	18.17	18.69	36.15
1	5	12.60	12.60	12.53	13.50	13.26	20.33
	10	F	20.78	19.89	18.30	18.87	36.15
2	5	12.13	12.08	12.21	13.49	13.26	20.33
	10	F	20.25	20.28	18.30	18.87	36.15
5	5	11.66	11.83	12.12	13.49	13.26	20.33
	10	F	F	20.24	18.30	18.87	36.15
10	5	12.19	12.24	12.11	13.49	13.26	20.33
	10	F	20.23	20.24	18.30	18.87	36.15

Table C.7: Results for scalability case when relaxing force. Iterations for the scheme depending on the number of previous time steps used (histories), iterations in each time step (ranking) and filter ( $\epsilon$ ) when relaxing displacement. F meaning a divergent experiment. The average number of iterations is 18.09 (sd=6.37).

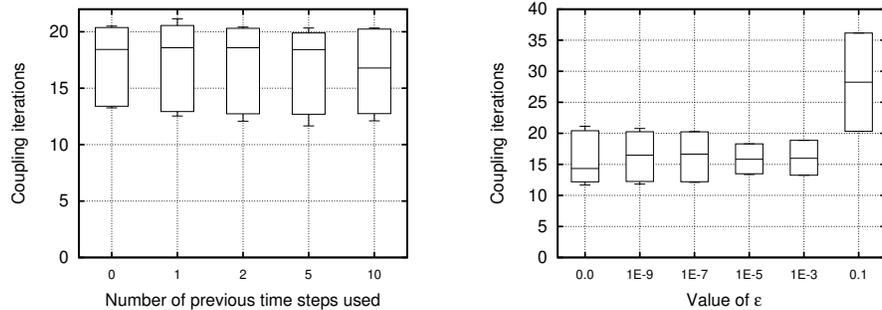


Figure C.13: Candle plots for the scalability case when force is relaxed.

## References

- [1] S. Deparis, M. Discacciati, G. Fourestey, A. Quarteroni, Fluid-structure algorithms based on Steklov-Poincaré operators, *Computer Methods in Applied Mechanics and Engineering* 195 (2006) 5797–5812.

- [2] M. W. Gee, U. Kuttler, W. A. Wall, Truly monolithic algebraicmultigrid for fluidstructure interaction, *International Journal for Numerical Methods in Biomedical Engineering* 85 (2010) 987–1016.
- [3] P. Crosetto, P. Reymond, S. Deparis, D. Kontaxakis, N. Stergiopoulos, A. Quarteroni, Fluid structure interaction simulation of aortic blood flow, *Computers and Fluids* 43 (2011) 46–57.
- [4] J. Hron, S. Turek, A monolithic fem/multigrid solver for an ale formulation of fluid-structure interaction with applications in biomechanics, *Fluid-Structure Interaction* 53 (2006) 146–170.
- [5] J. Degroote, K.-j. Bathe, J. Vierendeels, Performance of a new partitioned procedure versus a monolithic procedure in fluid structure interaction, *Computers and Structures* 87 (2009) 793–801.
- [6] H. G. Matthies, J. Steindorf, Partitioned strong coupling algorithms for fluid-structure interaction, *Computers and Structures* 81 (2003) 805–812.
- [7] C. Habchi, S. Russeil, D. Bougeard, J. L. Harion, T. Lemenand, A. Ghanem, D. D. Valle, H. Peerhossaini, Partitioned solver for strongly coupled fluid-structure interaction, *Computers and Fluids* 71 (2013) 306–319.
- [8] L. Radtke, A. Larena-Avellaneda, E. S. Debus, A. Düster, Convergence acceleration for partitioned simulations of the fluid-structure interaction in arteries, *Computational Mechanics* 57 (2016) 901–920.
- [9] F. Verdugo, W. A. Wall, Unified computational framework for the efficient solution of n-field coupled problems with monolithic schemes, *Computer Methods in Applied Mechanics and Engineering* 310 (2016) 335–366.
- [10] S. Badia, A. Quaini, A. Quarteroni, Modular vs. non-modular preconditioners for fluid-structure systems with large added-mass effect, *Computer Methods in Applied Mechanics and Engineering* 197 (2008) 4216–4232.

- [11] G. Houzeaux, M. Vázquez, R. Aubry, J. M. Cela, A massively parallel fractional step solver for incompressible flows, *Journal of Computational Physics* 228 (2009) 6316–6332.
- [12] G. Houzeaux, R. Aubry, M. Vázquez, Extension of fractional step techniques for incompressible flows: The preconditioned Orthomin(1) for the pressure Schur complement, *Computers and Fluids* 44 (2011) 297–313.
- [13] E. Casoni, A. Jérusalem, C. Samaniego, B. Eguzkitza, P. Lafortune, D. D. Tjahjanto, X. Sáez, G. Houzeaux, M. Vázquez, Alya: Computational Solid Mechanics for Supercomputers, *Archives of Computational Methods in Engineering* 22 (2015) 557–576.
- [14] M. Vazquez, G. Houzeaux, S. Koric, A. Artigues, J. Aguado-Sierra, R. Aris, D. Mira, H. Calmet, F. Cucchietti, H. Owen, A. Taha, J. M. Cela, Alya: Towards Exascale for Engineering Simulation Codes (2014) 1–20.
- [15] C. Förster, W. A. Wall, E. Ramm, Artificial added mass instabilities in sequential staggered coupling of nonlinear structures and incompressible viscous flows, *Computer Methods in Applied Mechanics and Engineering* 196 (2007) 1278–1293.
- [16] P. Causin, J. F. Gerbeau, F. Nobile, Added-mass effect in the design of partitioned algorithms for fluid-structure problems, *Computer Methods in Applied Mechanics and Engineering* 194 (2005) 4506–4527.
- [17] H. J. Bungartz, F. Lindner, M. Mehl, B. Uekermann, A plug-and-play coupling approach for parallel multi-field simulations, *Computational Mechanics* 55 (2015) 1119–1129.
- [18] B. W. Uekermann, Partitioned fluid-structure interaction on massively parallel systems, Ph.D. thesis, Technische Universität München, 2016.
- [19] R. Haelterman, A. E. Bogaers, K. Scheufele, B. Uekermann, M. Mehl, Improving the performance of the partitioned QN-ILS procedure for fluid-

- structure interaction problems: Filtering, *Computers and Structures* 171 (2016) 9–17.
- [20] M. Mehl, B. Uekermann, H. Bijl, D. Blom, B. Gatzhammer, A. Van Zuijlen, Parallel coupling numerics for partitioned fluid-structure interaction simulations, *Computers and Mathematics with Applications* 71 (2016) 869–891.
- [21] J. Degroote, Partitioned Simulation of Fluid-Structure Interaction: Coupling Black-Box Solvers with Quasi-Newton Techniques 20 (2013) 185–238.
- [22] K. Scheufele, Robust Quasi-Newton Methods for Partitioned Fluid-Structure Simulations, Ph.D. thesis, 2015.
- [23] J. Vierendeels, L. Lanoye, J. Degroote, P. Verdonck, Implicit coupling of partitioned fluid-structure interaction problems with reduced order models, *Computers and Structures* 85 (2007) 970–976.
- [24] A. E. Bogaers, S. Kok, B. D. Reddy, T. Franz, Quasi-Newton methods for implicit black-box FSI coupling, *Computer Methods in Applied Mechanics and Engineering* 279 (2014) 113–132.
- [25] G. Houzeaux, R. Borrell, J. Cajas, M. Vzquez, Extension of the parallel sparse matrix vector product (spmv) for the implicit coupling of pdes on non-matching meshes, *Computers and Fluids* 173 (2018) 216 – 225.
- [26] R. Borrell, J. Cajas, D. Mira, A. Taha, S. Koric, M. Vzquez, G. Houzeaux, Parallel mesh partitioning based on space filling curves, *Computers and Fluids* 173 (2018) 264 – 272.
- [27] G. Houzeaux, J. Principe, A variational subgrid scale model for transient incompressible flows, *International Journal of Computational Fluid Dynamics* 22 (2008) 135–152.
- [28] R. Calderer, A. Masud, A multiscale stabilized ALE formulation for incompressible flows with moving boundaries, *Computational Mechanics* 46 (2010) 185–197.

- [29] S. Turek, J. Hron, Proposal for Numerical Benchmarking of Fluid-Structure Interaction between an Elastic Object and Laminar Incompressible Flow, *Fluid-Structure Interaction (2006)* 371–385.
- [30] J. C. Cajas, M. Zavala, G. Houzeaux, E. Casoni, M. Vázquez, C. Moulinec, Y. Fournier, Fluid structure interaction in HPC multi-code coupling, in: *Civil-Comp Proceedings*, volume 107, 2015, pp. 1–26.
- [31] J. Degroote, P. Bruggeman, R. Haelterman, J. Vierendeels, Stability of a coupling technique for partitioned solvers in FSI applications, *Computers and Structures* 86 (2008) 2224–2234.
- [32] A. Santiago, J. Aguado-Sierra, M. Zavala-Aké, R. Doste-Beltran, S. Gómez, R. Arís, J. C. Cajas, E. Casoni, M. Vázquez, Fully coupled fluid-electromechanical model of the human heart for supercomputers, *International Journal for Numerical Methods in Biomedical Engineering* 34 (2018).