The 9th European Congress on Computational Methods in Applied Sciences and Engineering
ECCOMAS Congress 2024
3–7 June 2024, Lisboa, Portugal

A NEW APPROACH TOWARDS ASSESSING CONVERGENCE OF PARTITIONED FLUID-STRUCTURE INTERACTION SIMULATIONS

NICOLAS DELAISSÉ[1], THOMAS SPENKE[2], NORBERT HOSTERS[2] AND JORIS DEGROOTE[1,3]

[1]Department of Electromechanical, Systems and Metal Engineering
Ghent University
Sint-Pietersnieuwstraat 41, 9000 Ghent
e-mail: Belgium, Nicolas.Delaisse@UGent.be

[2]Chair for Computational Analysis of Technical Systems (CATS)
RWTH Aachen University
Schinkelstraße 2, Aachen, Germany

[3]Flanders Make @ UGent – Core Lab MIRO

Key words: Convergence Criterion, Partitioned Algorithm, Fluid-Structure Interaction, Coupled Problems

Summary. A partitioned algorithm couples distinct solvers by exchanging data on their common interface. In the case of strong coupling, where multiple coupling iterations are performed per time step, a convergence criterion is required to indicate that a time step has converged. A variety of such criteria exists, but they typically compare a quantity representing the change in interface data between successive coupling iterations, e.g., the interface residual, to a prescribed tolerance. The choice of its value is often ad hoc and based on experience. Moreover, its relationship to the tolerances of the subproblems is unclear and often not considered, even though the accuracy of the coupled simulation is inherently linked to that of the subproblem solutions.

This work addresses this shortcoming by discussing and investigating a newly introduced criterion that does not require the choice of a coupling tolerance, but bases the convergence of the time step on the number of subproblem iterations required to reach convergence. The new criterion can be applied to black-box solvers under the condition that they provide information on whether they have converged and on how many subproblem iterations have been run, which is not a very restrictive assumption in practice. Additionally, the solvers must be able to start from the solution of the previous coupling iteration. Although this work focuses on fluid-structure interaction simulations, the criterion is also applicable to other multiphysics problems.

1 INTRODUCTION

Partitioned algorithms for fluid-structure interaction (FSI) couple separate flow and structural solvers by exchanging data on their common interface. The goal is to ensure that the displacement of this interface is identical in both solvers, and that the forces on this interface are equal in size but opposite in sign. These two requirements are the kinematic and dynamic equilibrium conditions, respectively.

In strongly (implicitly) coupled problems, these conditions are satisfied by evaluating the solvers multiple times per time step, each time updating the interface boundary conditions of one solver based on interface data from the other. A single evaluation of both the flow and structural solver is called a *coupling iteration*. The degree to which the equilibrium conditions are fulfilled is commonly determined by comparing a norm of a change of interface data between consecutive coupling iterations to a so-called *coupling tolerance*. Only when this change drops below the prescribed value, the time step is considered converged and the next one can be started.

There are many variants for calculating a single scalar measure for the change in interface data, but for the most part they make use of the *coupling residual*. Focusing on the Gauss-Seidel type, where the flow and structural solver are called sequentially, this quantity can be defined as the difference between the output interface data of the structural solver and the corresponding prior input interface data of the flow solver. As a consequence, for Dirichlet-Neumann decomposition, this coupling residual is expressed in units of displacement. A similar value is sometimes defined for the traction forces which are the out- and input of the flow and structural solver, respectively.

The most straightforward approach is to compare the Euclidean norm of the coupling residual to a tolerance value [1–5]. The disadvantage of this approach is that the value of the residual will depend on the discretization. To alleviate this, the norm can be scaled with the square root of the number of degrees of freedom on the interface [6, 7]. Still, the residual will depend on the magnitude of the displacement values. Therefore, the absolute norm can be normalized using either the initial residual norm at the start of the time step [2, 8] or the Euclidean norm of the interface data itself [4, 9, 10]. Finally, a combination of these criteria can be applied, and while often only the displacement coupling residual is checked, a criterion can also be defined for the force interface residual.

The common factor among the multitude of criteria is that they require to specify a tolerance. The choice of an appropriate value is not straightforward and often made based on experience or trial-and-error. Moreover, its value does not independently determine the performance of the coupled simulation in terms of accuracy and run time but is also strongly linked to the tolerances of the solvers themselves. For example, if the solver tolerances are set very strict while the coupling tolerance is loose, a lot of time is invested in accurate solving the separate domains without it being reflected in the quality of the coupled solution because the equilibrium conditions are only loosely satisfied. Conversely, if the coupling tolerance is too strict compared to the solver tolerances, convergence may never be achieved. This is because the decrease of the coupling residual and thus the satisfaction of the equilibrium conditions is limited by the numerical solution errors of the separate solvers. In other words, the remaining change in interface data between coupling iterations results from the solution errors of the solvers and might not be reduced by performing more coupling iterations. Moreover, also the solver tolerances are linked to each other. For instance, if the tolerance of the structural solver is much stricter than that of the flow solver, computational time will be misspend on resolving the structural domain. Obviously, the inverse is also true.

The choice of a coupling tolerance and by extension the solver tolerances is very problem-dependent. First of all, the convergence criteria for the solvers differ considerably between solver packages. While they commonly compare one or more integral residual measures to prescribed tolerances, a direct comparison with the tolerances of other solver codes or the coupling tolerance

is typically not possible. Second, appropriate values for the tolerances depend on the physics occurring as well as the size, geometry and discretization of the domain. For example, some problems might feature large displacement and small pressure differences, while for others the inverse could be true. Third, in most problems the solution of one of the subsolvers requires more effort than the other, which also affects the optimal choice.

The main question is how to assess which choices lead to the desired accuracy without overspending computational resources. As part of the solution, this work discusses a new coupling criterion [11] that avoids the need of a coupling tolerance but bases the convergence of the time step on the number of iterations each of the subproblem requires to reach convergence. Although the focus of this work is on fluid-structure interaction problems, the concepts and results can be translated to any type of partitioned coupling of black-box solvers.

The next section introduces the necessary notation and a typical convergence criterion using a coupling tolerance. Thereafter, Section 3 discusses the new coupling criterion and provides details on its implementation. Section 4 presents simulation results for two test cases: the flexible tube and lid-driven cavity. Finally the main conclusions are presented in Section 5.

2 PARTITIONED FLUID-STRUCTURE INTERACTION

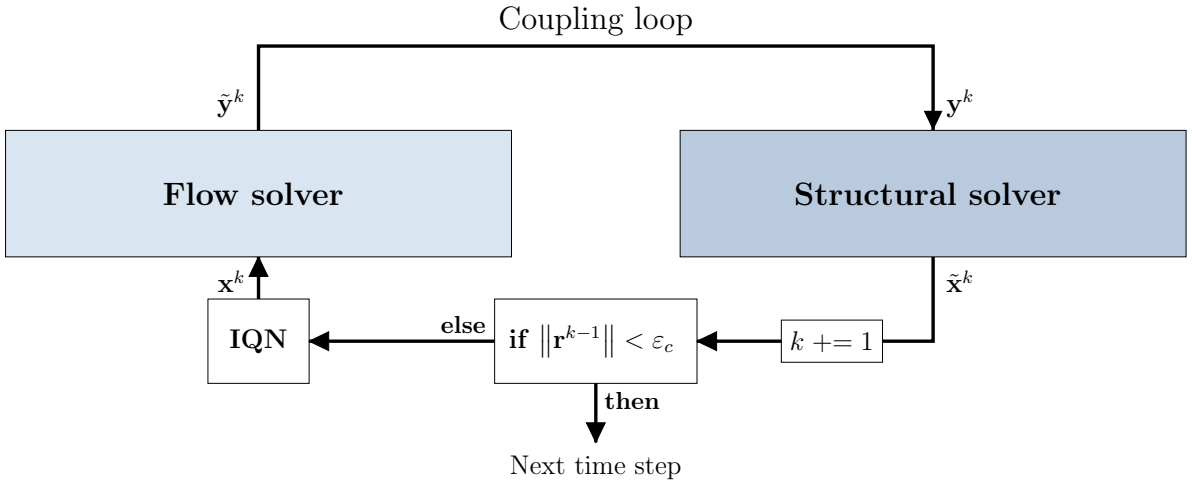

Figure 1: Illustration of the coupling of the flow and structural solver, where the non-normalized criterion is used to determine convergence of the coupling. Each solver block is run until the corresponding solver tolerance is achieved.

Fig. 1 illustrates a typical setup for partitioned coupling of a flow and structural solver, where in each call the solvers are run to full convergence. The displacement data on the interface is represented by the vector $\mathbf{x}$ (unit m), while the pressure and traction forces are grouped in the vector $\mathbf{y}$ (unit Pa). To avoid divergence and accelerate convergence of strongly coupled problems, it is often required to update the interface data between solvers, through, for example, relaxation or the use a coupling algorithm such as the interface quasi-Newton (IQN) type [2]. Therefore, the output of the solvers is distinguished by the use of a tilde. Generally, the interface displacement and forces will not be identically discretized in both solvers, necessitating spatial interpolation. Since this is not the focus of the current work, it is tacitly assumed to occur.

With these notations, the kinematic and dynamic equilibrium conditions can be expressed as

$$\left\|\tilde{\mathbf{x}}^k - \mathbf{x}^k\right\| < \varepsilon_x \quad \text{and} \quad \left\|\tilde{\mathbf{y}}^k - \mathbf{y}^{k-1}\right\| < \varepsilon_y, \tag{1}$$

where the subscript k is denotes the coupling iteration ($k = 0, 1, \ldots$ in every time step).

Often, only the displacement interface data is modified between the solvers, for example by IQN, and only the kinematic equilibrium condition is checked. The difference $\tilde{\mathbf{x}}^k - \mathbf{x}^k$ is then regarded as the coupling residual with unit m, defined as $\mathbf{r}^k = \tilde{\mathbf{x}}^k - \mathbf{x}^k$.

As detailed in Section 1, the convergence is typically assessed by requiring that the norm of this coupling residual is sufficiently small. In this work, the non-normalized convergence criterion will be used as the typical convergence criterion to be compared to the new criterion explained in the next session. This convergence criterion is expressed as

$$\left\|\mathbf{r}^k\right\| < \varepsilon_c, \tag{2}$$

where ε_c is the coupling tolerance with unit m.

3 NEW COUPLING CRITERION

3.1 Motivation

In Spenke and Delaissé et al. [11], a new coupling criterion is introduced that bases the convergence of the coupling loop on the convergence of the solvers and avoids the use of a coupling tolerance. A closer look at the solvers is needed to justify this criterion.

A great advantage of the partitioned coupling approach is that solvers can be black boxes meaning that no access to the solver codes is required. Nevertheless, since the flow and structural problems are typically nonlinear, each solver can be assumed to have an internal iteration loop, as shown in Fig. 2. In these internal loops, so-called *subproblem iterations* ($i = 0, 1, \ldots$ in every solver call) are performed until the *subproblem residual* $\mathbf{r}_p^i$ becomes smaller than a prescribed *solver tolerance* ε_p, where the subscript p is a general notation for f or s, respectively referring to the flow and structural subproblem. The exact implementation differs strongly between solution methods and implementations, but in general the loop consists of three consecutive steps: assembling the system, evaluating the residual and solving the system.

Here, it is assumed that all dependency on the solution field $\mathbf{u}$ inherent to the subproblem is included in the system matrix $\mathbf{A}(\mathbf{u})$. Therefore, the right-hand side (RHS) $\mathbf{b}$ containing the boundary conditions and thus also the interface data received from the other solver stays constant within each subproblem call. When coupled to another subproblem as in partitioned FSI, the received interface data contained by the RHS has been determined by the other solver using the solution field $\mathbf{u}$. As a consequence, the RHS of the solver is also dependent on $\mathbf{u}$, even if only in an implicit manner. For example, a change in position of the interface will affect the fluid forces and hence also the boundary condition for the structural subproblem. Analogously, the structural solver uses the forces on the interface to calculate the interface position which, in turn, influences the flow solution and thus the pressure and traction forces on the interface.

The received interface data and RHS $\mathbf{b}$ are determined based on the last subproblem solution, which is also the initial iterate $\mathbf{u}^0$ for the next subproblem evaluation. Hence, the RHS can be written as $\mathbf{b}|_{\mathbf{u}^0}$, where the vertical line expresses the implicit dependence through the other subproblem.

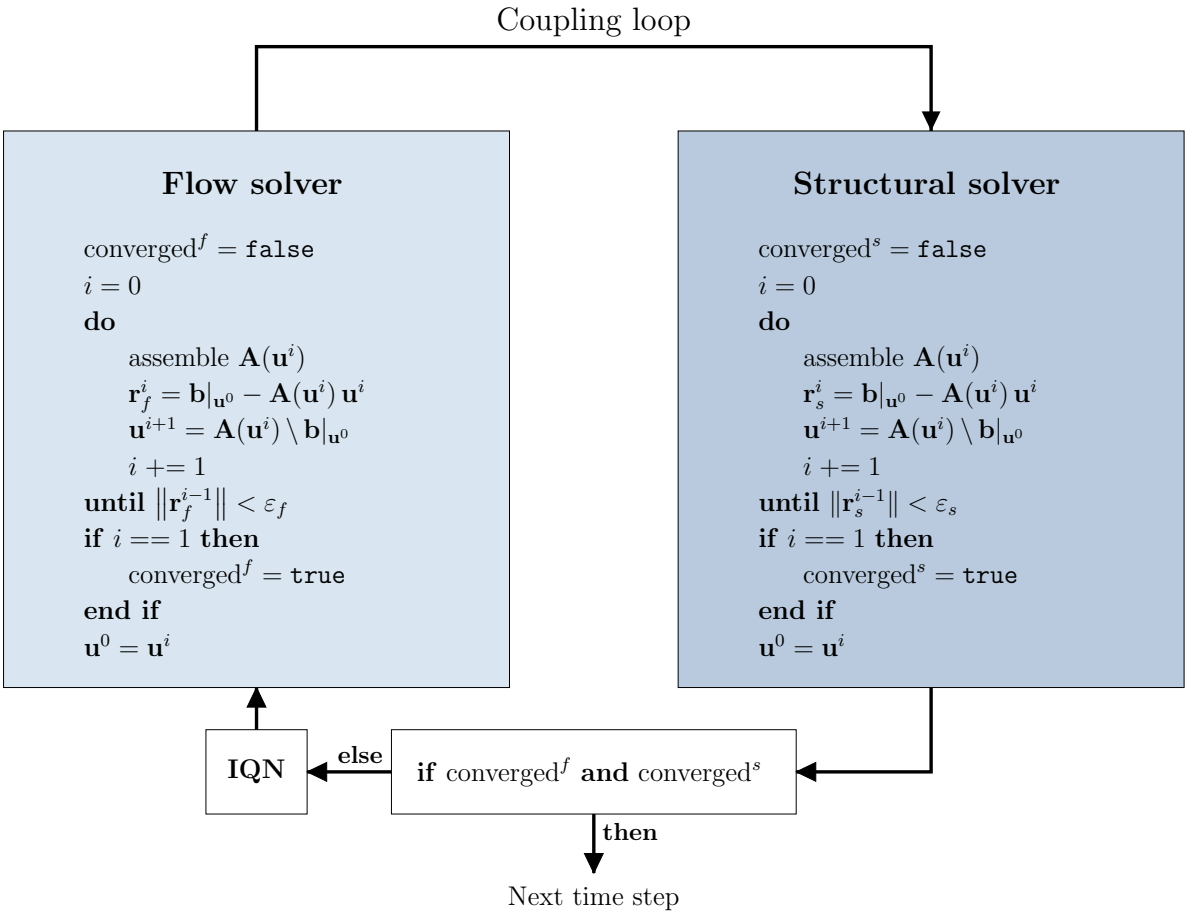

Figure 2: Illustration of the new coupling criterion [11] with three iteration loops: the coupling loop and an internal iteration loop in each subproblem. Note that the assembly is done with the previous solution. The notation $\mathbf{A}(\mathbf{u}^i) \setminus \mathbf{b}|_{\mathbf{u}^0}$ is symbolic and refers to solving the system. Similarly, $\|\mathbf{r}_p^{i-1}\| < \varepsilon_p$ $(p = f, s)$ is a notation for the convergence criterion of the solvers.

Considering the subproblem residual $\mathbf{r}_p^i$ in subproblem iteration i

$$\mathbf{r}_p^i = \mathbf{b}|_{\mathbf{u}^0} - \mathbf{A}(\mathbf{u}^i)\,\mathbf{u}^i \quad \text{for} \quad i = 0, 1, \ldots, \tag{3}$$

it becomes clear that only the initial subproblem residual of the solver call $\mathbf{r}_p^0$ considers left- and right-hand sides referring to the same solution field $\mathbf{u}^0$. For the subsequent subproblem iterations ($i > 0$), the left-hand side $\mathbf{A}(\mathbf{u}^i)\,\mathbf{u}^i$ is updated, while the RHS stays unchanged and refers to the outdated $\mathbf{u}^0$. The updated RHS $\mathbf{b}|_{\mathbf{u}^i}$ would only become available after solving the other subproblem using $\mathbf{u}^i$ as boundary condition.

This finding motivates the use of $\mathbf{r}_f^0$ and $\mathbf{r}_s^0$ to assess the convergence of the coupling loop: the coupling is considered converged when both solvers fulfill their convergence criteria in the first subproblem iteration:

$$\|\mathbf{r}_f^0\| < \varepsilon_f \quad \text{and} \quad \|\mathbf{r}_s^0\| < \varepsilon_s. \tag{4}$$

In other words, the change in the RHS of the subproblems between two coupling iterations has to be so small that the solvers immediately converge.

An additional benefit of the new criterion is that it is not required to fully converge the solvers in each call. Instead, a maximal number of subproblem iterations $n_{max,p}$ ($p = f, s$) can be defined for each solver. It has been shown that this can reduce the run time of the partitioned simulation, while maintaining the accuracy [11]. Doing so with the typical convergence criteria is not possible without additional checks, since there is no guarantee that the solvers are converged once the coupling converges.

In Fig. 2, the residual is determined before the system is solved. This means that one more update is performed even when the subproblem has converged. While this is not strictly necessary, it is often impossible to change in black-box solvers. Moreover, by omitting this update, the accuracy of the coupled simulation would be determined by the tolerances that is satisfied first, rather than the one that is satisfied last.

The reasoning in this section in principle only holds when the interface data is not changed between the solvers, while in practice, this is often done as part of relaxation or a quasi-Newton update. However, the update increment of the interface data is proportional to the interface residual and therefore vanishes upon convergence as is also confirmed in [11]. The effect of the update can therefore safely be neglected.

3.2 Implementation

To implement the new criterion, it is often simplest to check after each solver call whether the solver has converged and whether only one subproblem iteration was run. Alternatively, one could start with a solver call allowing only one subproblem iteration and check its convergence. If the solver did not converge, another solver call is made with a maximal number of subproblem iterations $n_{max,p} - 1$. The latter approach was followed for the two solvers used in this work: OpenFOAM 11 [12] for the fluid and the Sturctural Mechanics Application of Kratos Multiphysics 9.4 [13] for the structure. The Python code CoCoNuT (GitHub repository pyfsi/coconut) establishes the coupling and the application of the new criterion.

Since the solvers are considered black boxes, their solution procedures are not modified. For instance, OpenFOAM is a finite-volume solver that uses a segregated pressure-velocity coupling, meaning it sequentially solves for velocity and pressure. Therefore, a flow solver call is considered converged if both the initial residual of the velocity equation and the initial residual of the pressure equation calculated with the updated velocity are less than ε_f. Once this occurs, the solver performs a final iteration, where relaxation of the equations can be turned off. In this work, no relaxation factors are applied in order not to obscure the effects of interest, but the final flow solver iteration is still run. As a consequence, when the initial residual converges, still two subproblem iterations are performed.

4 RESULTS AND DISCUSSION

4.1 Test cases

Two test cases are considered. The first is a three-dimensional flexible tube enclosing an incompressible fluid [14], shown in Fig. 3. In this case, a pressure pulse at the inlet causes a wave to travel through the tube at a finite velocity. The second is a two-dimensional cavity with a flexible bottom [15, 16], shown in Fig. 4. By applying a periodically varying horizontal velocity at the top, the flexible membrane starts to oscillate.

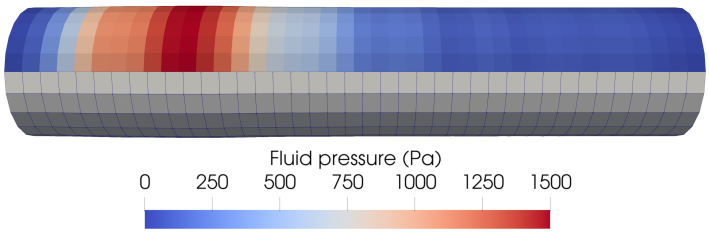

Figure 3: Illustration of the flexible tube at $t = 0.004\,\mathrm{s}$. The upper half shows the fluid pressure field in Pa and the lower part illustrates shell elements of the structure.

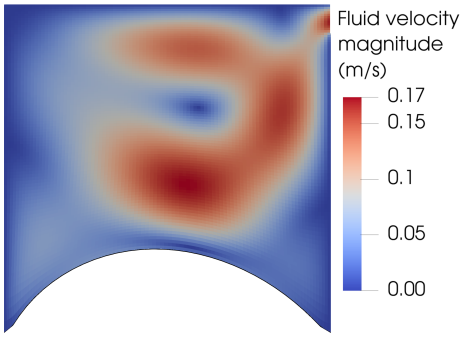

Figure 4: Illustration of the lid-driven cavity case at $t = 65\,\mathrm{s}$, showing the deformation of the flexible bottom and the corresponding velocity field in m/s.

The material parameters and boundary conditions are adopted from literature [1, 15]. The discretizations of the subproblems are identical to those of the FV-FE framework described in Spenke and Delaissé et al. [11], except for the flow domain of the lid-driven cavity which has been refined with a factor three in both the vertical and horizontal direction. Here OpenFOAM is used as the flow solver instead of Ansys Fluent. Because of the segregated pressure-velocity coupling in OpenFOAM, the time step of the lid-driven cavity case had to be decreased with a factor ten to 0.01 s. To limit the computational time of the study presented hereafter, the lid-driven cavity case is always restarted from a saved state at $t = 60\,\mathrm{s}$ and run up to $t = 70\,\mathrm{s}$.

The coupling algorithm for these FSI problems is IQN-ILSM [17] with reuse as surrogate model. For the flexible tube case, the last ten time steps are reused ($q = 10$) and none ($q = 0$) for the lid-driven cavity case. No filtering of linear dependencies in the residual data is applied, as this is not really required for IQN-ILSM, as opposed to IQN-ILS where it is important for cases with reuse.

4.2 Parameter study

To gain insight in the effects of the solver and coupling tolerances as well as to study the performance of the new criterion, a parameter study is performed for each test case, visualized in Figs. 5 and 6. Two tolerance levels are selected for each solver, resulting in four combinations of flow and structural solver tolerances, each indicated by its own color. Several different coupling criteria are examined:

- First, a series of coupling tolerances ε_c are chosen which are used to determine convergence of the coupling loop as explained in Section 2.
- Further, another set of runs uses the new criterion.
- In order to compare these two criteria, a third type of runs is performed, for which the coupling tolerance ε_c has been chosen in such a way that the norm of the coupling residual at the end of each time step is approximately equal to that of the new criterion. This will allow to compare the performance in terms of accuracy and time. Since these runs are not obtainable without running the new criterion first, their colors are faded.

- Next, another set of runs is performed where the new criterion is applied but now also limiting the number of subproblem iterations.
- Finally, a reference run (purple) is performed with strict solver and tolerance settings which will allow to evaluate the error on the interface, as will be detailed below.

For each bar in these plots, the run was performed three times and the mean is reported. The error bars in black representing one standard error of the mean provide a measure for the spread.

4.2.1 Flexible tube

Fig. 5 gives an overview of the parameter study for the flexible tube case. Fig. 5a shows the average norm of the coupling residual $\|\mathbf{r}\|$ at the end of each time step. This value is lower than the prescribed tolerance ε_c which sets the upper limit for each time step. Here, the obtained value is each time about half of the prescribed value. The coupling tolerances $\varepsilon_c = 10^{-10}\,\mathrm{m}$ and $\varepsilon_c = 10^{-3}\,\mathrm{m}$ are not shown because the simulation diverges for all four solver tolerance combinations. For the reference run, a coupling tolerance of $\varepsilon_c = 10^{-13}\,\mathrm{m}$ is set and the solver tolerances for the flow and structure are $\varepsilon_f = 10^{-8}$ and $\varepsilon_s = 10^{-12}$. Also the tolerances for the solution of the subsystems of the flow solver (e.g., velocity and pressure equations) have been decreased. The figure clearly illustrates that with the new criterion, the final coupling residual is determined by the choice of solver tolerances, independent of the number of subproblem iterations per solver call.

In Fig. 5b, the effect on calculation time is reported. As expected, a stricter coupling tolerance generally leads to a higher computational cost. However, for the loosest tolerance ($\varepsilon_c = 10^{-4}$), the run time increases for the two runs with the strict structural solver tolerance (orange and red). In that regard, it is important to note that the structural part of the calculation for this test case takes about twice as much time as the flow part, and therefore, the value of the structural tolerance will have a larger effect. Comparing the new criterion to the use of a coupling tolerance (faded color) 11235813 reveals that the efficiency in terms of run time is comparable, but the new criterion is on average a 8 % slower. However, if the number of subproblem iterations is limited, the new criterion is on average 25 % faster for the runs with the strictest structural solver tolerance (orange and red).

Furthermore, Figs. 5c and 5d show the accuracy in terms of both interface displacement $\mathbf{x}$ and interface forces $\mathbf{y}$. This error is calculated by comparing the interface solution with the reference run. The average norm of these difference is divided by the number of degrees of freedom on the interface n to result in a root-mean-square error (RMSE). For example, for the displacement, this value is given by

$$\mathrm{RMSE}_{\mathbf{x}} = \frac{1}{m} \sum_{m} \frac{\|\mathbf{x} - \mathbf{x}_{ref}\|}{\sqrt{n}}, \tag{5}$$

where m is the number of time steps. The trends and the number of orders of magnitudes that are spanned in these figures are very close to those in Fig. 5a.

Figs. 5e and 5f visualize the trade-off between accuracy and computational cost by plotting all runs with respect to these criteria. The markers with a black outline are located on the *Pareto front* [18], which means that they are optimal in the sense that there are no runs better in

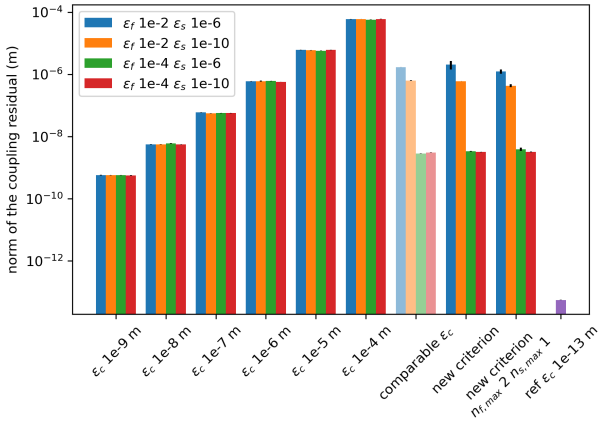

(a) Norm of the coupling residual.

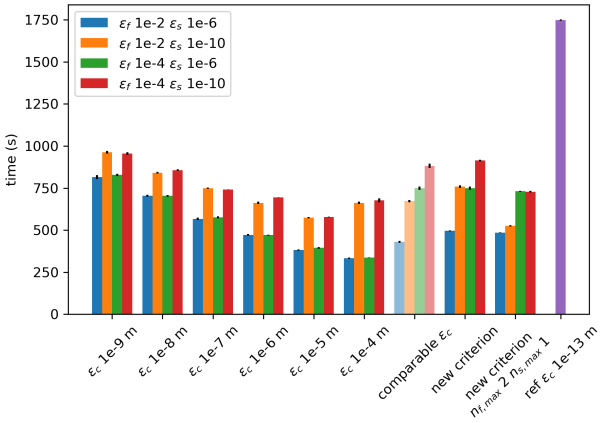

(b) Run time.

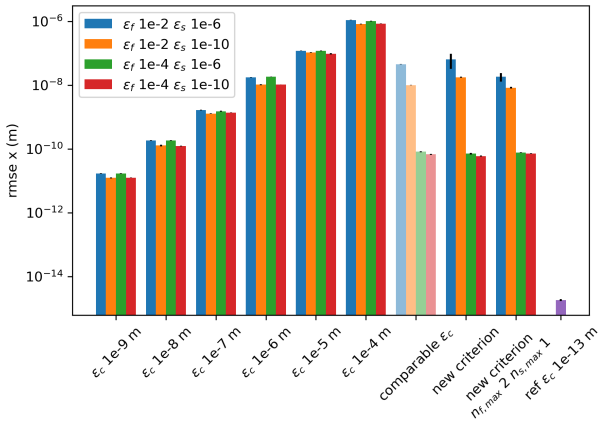

(c) RMSE of the interface displacement.

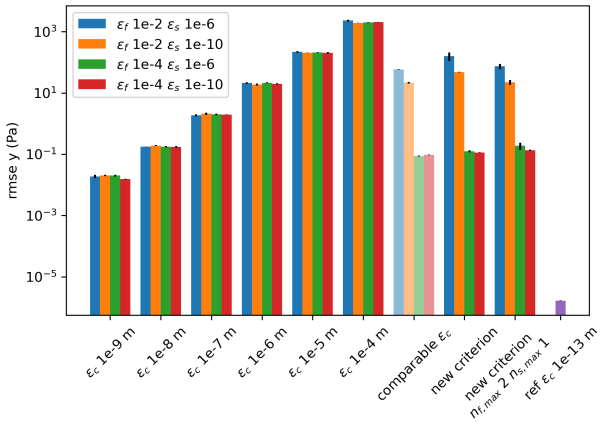

(d) RMSE of the interface forces.

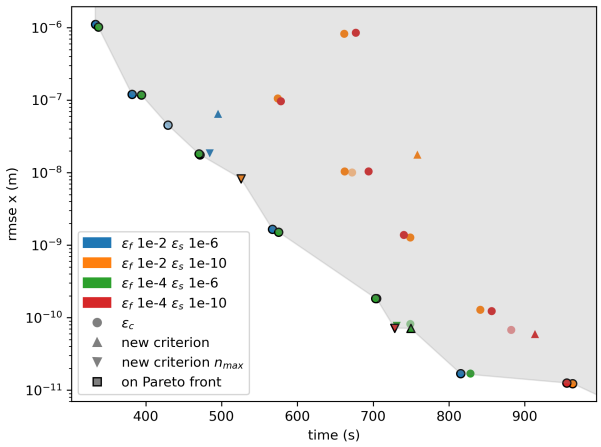

(e) Pareto front of displacement error and time.

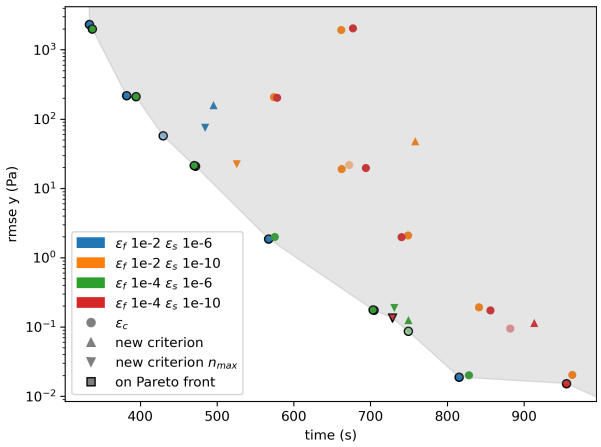

(f) Pareto front of forces error and time.

Figure 5: Parameter study for the flexible tube case. The error bars in black represent one standard error of the mean. In e and f, the area indicated in gray is the dominated region bounded by the Pareto front and the faded colors refer to the runs with a coupling tolerance chosen to match the new criterion.

terms of both accuracy and computational cost. Conversely, for the other so-called dominated runs, another run exists that performs better in both criteria. Clearly, the runs with strict structural tolerance perform worse. Nevertheless, with new criterion and by limiting the number of subproblem iterations, the runs with these solver settings are on or close to the Pareto front.

4.2.2 Lid-driven cavity

In Fig. 6, the parameter study for the lid-driven cavity case is presented. Figs. 6a and 6b show similar trends as the flexible tube case, but the optimal coupling tolerance is now $\varepsilon_c = 10^{-6}\,\mathrm{m}$, which is the second most strict coupling tolerance for which the calculation converged. For the coupling tolerance $\varepsilon_c = 10^{-7}\,\mathrm{m}$, the runs did not consistently converge. Here, the new criterion does not lead to a speed-up, not even if the number of subproblem iterations is limited. The remaining advantage of the new criterion is that the user has to set one fewer tolerance.

By comparing with a reference run, root-mean-square errors are again calculated with Eq. (5). This reference run uses solver tolerances for the flow and structure equal to $\varepsilon_f = 10^{-6}$ and $\varepsilon_s = 10^{-12}$. While the behavior of the $\mathrm{RMSE}_{\mathbf{y}}$ (Fig. 6d) is once more very similar to the coupling residual, this is not the case for $\mathrm{RMSE}_{\mathbf{x}}$ (Fig. 6c).

Furthermore, Figs. 6e and 6f plot the trade-off between the computational cost and the accuracy. For this test case, there is no clear set of solver tolerances that performs best, but as for the flexible tube case, the runs with new convergence criterion are on or close to the Pareto front.

4.2.3 Remarks

The new criterion, Eq. (4), poses a condition for each solver. The solver tolerances can be balanced by choosing them such that both conditions are met in approximately the same coupling iteration (*orange* color). However, this did not lead to visible performance improvements.

The number of subproblem iterations can be optimized to achieve minimal run time for each set of solver tolerances when using the new criterion [11]. While doing so might result in a larger speed-up, this is postponed to future work.

5 CONCLUSIONS

Typically, partitioned FSI simulations require choosing three tolerances: besides the flow and structural tolerance needed for the subproblems, an additional tolerance is set for the coupling. The choice of this tolerance is not straightforward. Moreover, these three tolerances are not independent and have a significant impact on the accuracy and computational cost of the simulations.

In this work a new coupling criterion is discussed which avoids this additional tolerance and instead assesses the convergence of the coupling based on the convergence of the subproblems. This approach provides the additional benefit that the number of subproblem iterations can be limited which can reduce the run time.

With respect to performance, general conclusions are difficult to draw, since the relation between the tolerances are very case-dependent and depend on many factors, such as the computing infrastructure. However, for the two cases examined, the new criterion with limitation of the number of subproblem iterations resulted in (close to) optimal settings.

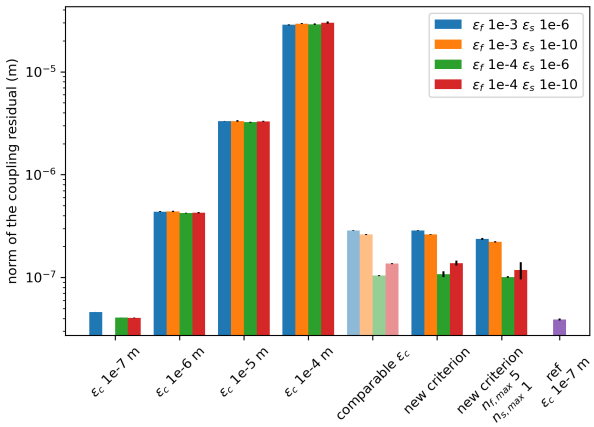

(a) Norm of the coupling residual.

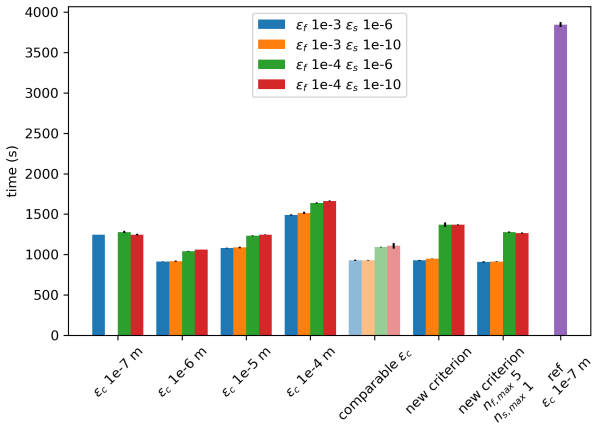

(b) Run time.

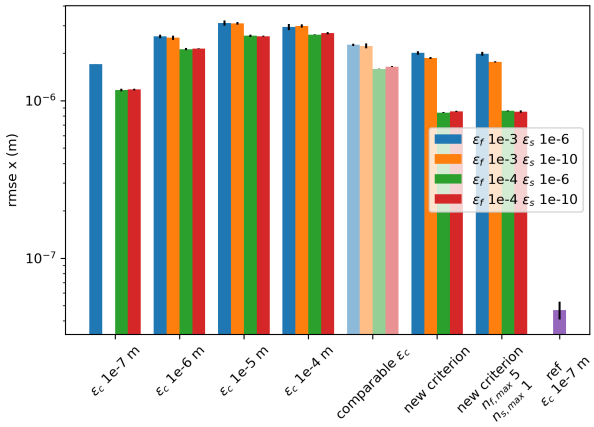

(c) RMSE of the interface displacement.

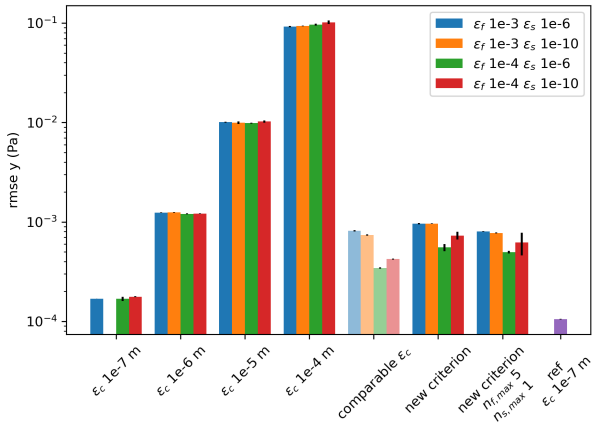

(d) RMSE of the interface forces.

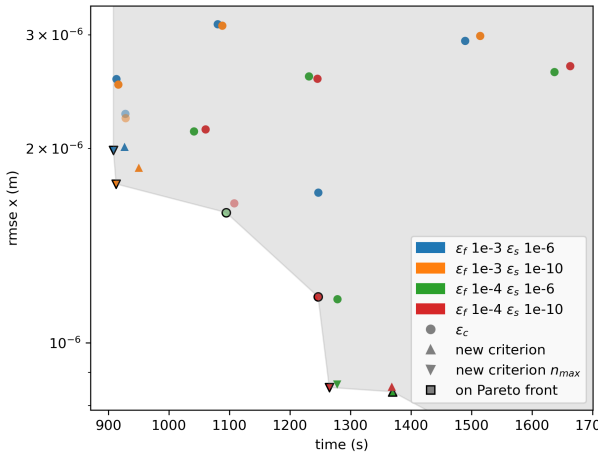

(e) Pareto front of displacement error and time.

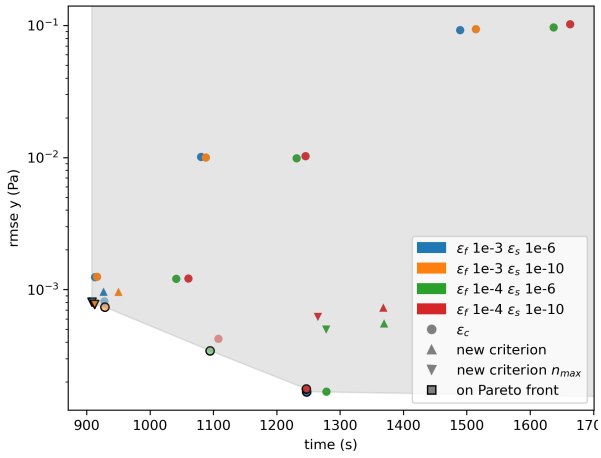

(f) Pareto front of forces error and time.

Figure 6: Parameter study for the lid-driven cavity case. The error bars in black represent one standard error of the mean. In e and f, the area indicated in gray is the dominated region bounded by the Pareto front and the faded colors refer to the runs with a coupling tolerance chosen to match the new criterion. For the coupling tolerance $\varepsilon_c = 10^{-7}$ m, some runs are missing since they did not consistently converge. Also the *red* run with limited number of subproblem iteration and the reference run diverged at times.

REFERENCES

[1] J. Degroote, R. Haelterman, S. Annerel, P. Bruggeman, and J. Vierendeels, "Performance of partitioned procedures in fluid-structure interaction," *Comput Struct*, vol. 88, no. 7–8, pp. 446–457, 2010. DOI: 10.1016/j.compstruc.2009.12.006.

[2] N. Delaissé, T. Demeester, R. Haelterman, and J. Degroote, "Quasi-Newton methods for partitioned simulation of fluid-structure interaction reviewed in the generalized Broyden framework," *Arch Comput Method E*, vol. 30, no. 5, pp. 3271–3300, Apr. 2023. DOI: 10.1007/s11831-023-09907-y.

[3] S. T. Ha and H. G. Choi, "Investigation on the effect of density ratio on the convergence behavior of partitioned method for fluid–structure interaction simulation," *J Fluid Struct*, vol. 96, p. 103 050, Jul. 2020. DOI: 10.1016/j.jfluidstructs.2020.103050.

[4] R. Haelterman, A. E. J. Bogaers, K Scheufele, B Uekermann, and M. Mehl, "Improving the performance of the partitioned QN-ILS procedure for fluid-structure interaction problems: Filtering," *Comput Struct*, vol. 171, pp. 9–17, 2016. DOI: 10.1016/j.compstruc.2016.04.001.

[5] T. Gallinger and K.-U. Bletzinger, "Comparison of algorithms for strongly coupled partitioned fluid-structure interaction - Efficiency versus simplicity," in *5th European Conference on Computational Fluid Dynamics*, J. C. F. Pereira and A. Sequeira, Eds., Lisbon, Portugal, Jun. 2010, pp. 1–20.

[6] R. Zorrilla and R. Rossi, "A memory-efficient multivector quasi-Newton method for black-box fluid-structure interaction coupling," *Comput Struct*, vol. 275, p. 106 934, Jan. 2023. DOI: 10.1016/j.compstruc.2022.106934.

[7] A. E. J. Bogaers, S. Kok, B. D. Reddy, and T. Franz, "Quasi-Newton methods for implicit black-box FSI coupling," *Comput Method Appl M*, vol. 279, pp. 113–132, 2014. DOI: 10.1016/j.cma.2014.06.033.

[8] R. Haelterman, J. Degroote, D. Van Heule, and J. Vierendeels, "The quasi-Newton least squares method: A new and fast secant method analyzed for linear systems," *Siam J Numer Anal*, vol. 47, no. 3, pp. 2347–2368, 2009. DOI: 10.1137/070710469.

[9] K. Davis, M. Schulte, and B. Uekermann, "Enhancing quasi-Newton acceleration for fluid-structure interaction," *Math Comput Appl*, vol. 27, no. 3, p. 40, May 2022. DOI: 10.3390/mca27030040.

[10] T. Spenke, N. Hosters, and M. Behr, "A multi-vector interface quasi-Newton method with linear complexity for partitioned fluid–structure interaction," *Comput Method Appl M*, vol. 361, p. 112 810, 2020. DOI: 10.1016/j.cma.2019.112810.

[11] T. Spenke, N. Delaissé, J. Degroote, and N. Hosters, "On the number of subproblem iterations per coupling step in partitioned fluid-structure interaction simulations," *Int J Numer Meth Eng*, vol. 125, no. 7, Jan. 2024. DOI: 10.1002/nme.7420.

[12] C. Greenshields, *OpenFOAM v11 User Guide*. London, UK: The OpenFOAM Foundation, 2023.

[13] P. Dadvand, R. Rossi, and E. Oñate, "An object-oriented environment for developing finite element codes for multi-disciplinary applications," *Arch Comput Method E*, vol. 17, no. 3, pp. 253–297, Jul. 2010. DOI: 10.1007/s11831-010-9045-2.

[14] J. Degroote, K.-J. Bathe, and J. Vierendeels, "Performance of a new partitioned procedure versus a monolithic procedure in fluid-structure interaction," *Comput Struct*, vol. 87, no. 11–12, pp. 793–801, Jun. 2009. DOI: 10.1016/j.compstruc.2008.11.013.

[15] D. P. Mok, "Partitionierte lösungsansätze in der strukturdynamik und der fluid-struktur-interaktion," de, Ph.D. dissertation, Universität Stuttgart, Stuttgart, Germany, 2001. DOI: 10.18419/OPUS-147.

[16] G. Valdés, "Nonlinear analysis of orthotropic membrane and shell structures including fluid-structure interaction," Ph.D. dissertation, Universitat Politècnica de Catalunya, Barcelona, Spain, 2007.

[17] N. Delaissé, T. Demeester, D. Fauconnier, and J. Degroote, "Surrogate-based acceleration of quasi-Newton techniques for fluid-structure interaction simulations," *Comput Struct*, vol. 260, p. 106 720, Feb. 2022. DOI: 10.1016/j.compstruc.2021.106720.

[18] P. Fleming, "Designing control systems with multiple objectives," in *IEE Master Class. Advances in Control Technology*, IEE, 1999. DOI: 10.1049/ic:19990716.