A two-scale approximation of the Schur complement and its use for non-intrusive coupling

Lionel Gendre, Olivier Allix, Pierre Gosselet

A two-scale approximation of the Schur complement and its use for non-intrusive coupling

L. Gendre, O. Allix, P. Gosselet
LMT-Cachan (ENS Cachan/CNRS/UPMC/PRES UniverSud Paris)
61 avenue du Président Wilson, F-94235 Cachan Cedex - FRANCE

June 11, 2020

Abstract

This paper presents a two-scale approximation of the Schur complement of a subdomain's stiffness matrix, obtained by combining local (*ie.* element strips) and global (*ie.* homogenized) contributions. This approximation is used in the context of a coupling strategy that is designed to embed local plasticity and geometric details into a small region of a large linear elastic structure; the strategy consists in creating a local model which contains the desired features of the concerned region and then substituting it into the global problem by the means of a non-intrusive solver coupling technique adapted from domain decomposition methods. Using the two-scale approximation of the Schur complement as a Robin condition on the local model enables to reach high efficiency. Examples include a large 3D problem provided by our industrial partner Snecma.

Keywords: Structures, Finite element methods, Multiscale analysis, Nonlinear solvers, Schur complement, Non-intrusive methods

1 INTRODUCTION

Substructuring and domain decomposition methods are very powerful analysis techniques in the field of structural mechanics. By splitting a large problem into several smaller subproblems, these techniques can help keeping computational costs at reasonable levels. However, their efficiency depends entirely on how well the subproblems are bridged together — that is, how well the interactions between each subdomain and the rest of the structure "that it sees" are described.

In the context of finite element analysis (FEA) applied to static linear elasticity problems, the "mechanical impedance" of a region can be described by the *Schur complement* of its stiffness matrix on its boundary. This quantity is involved in a broad range of techniques, from static condensation (which basically consists in computing it straightforwardly) to mixed domain decomposition methods [18, 4], and to the design of efficient preconditioners for primal or dual domain decomposition methods [16, 6, 19]. A classical technique for its computation is to perform a block Gaussian elimination with specific (and non-optimal) numbering. However, this technique is usually very expensive on large problems; for that reason, several approximations of the Schur complement have been investigated in the literature. Some of those techniques are reviewed in Section 3.

In this paper, a new approximation technique of the Schur complement is presented. Unlike most common approaches, which retain either short-scale effects (*ie.* the local influence of the vicinity of the interface) or long-scale effects (*ie.* the global influence of a large region around the interface), this technique consists in combining short-scale and long-scale information in order to obtain a realistic representation of a subdomain's stiffness at both scales, that is much cheaper to compute than the exact Schur complement. This approximation was designed in the framework of a non-intrusive coupling strategy that was introduced for the first time in [9]; this context explains some of the limitations and choices of the proposed Schur complement approximation. It should be noted, however, that the approximation technique we suggest is not restricted to this context, and can probably have many more applications, such as building preconditioners for domain decomposition methods.

The rest of the paper is structured as follows. In Section 2, we review the basics of the coupling strategy. Section 3 presents the application of the two-scale Schur complement approximation to the

strategy, as a way to build efficient local boundary conditions. Finally, Section 4 contains results obtained on a large 3D example provided by our industrial partner Snecma, as a proof of feasibility of the method.

2 A NON-INTRUSIVE COUPLING STRATEGY

2.1 Reference problem and splitted formulation

We consider a static mechanical problem set on a domain Ω, under small perturbation hypothesis. We assume that the behavior of the structure is linear elastic except in a small region denoted Ω_I, which is already known; in this region, we assume the constitutive law is elastic-plastic (although other forms of nonlinearity could as well be considered). The reference problem is shown on Figure 1(a) with its boundary conditions (schematized by the triangles) and loads (echematized by the arrows).

This problem can be rephrased in a splitted way which enables using a simplified global linear elastic model (Figure 1(b)), completed by an enhanced nonlinear local model (Figure 1(c)) which may contain geometric details (like holes) that do not appear in the global model. We assume the global model is analyzed using commercial (and closed) FEA software, whereas the local model can be analyzed using dedicated software. Our goal is to find the solution to the reference problem by using the global and local models and solvers in a *non-intrusive* way — that is, without modifying them.

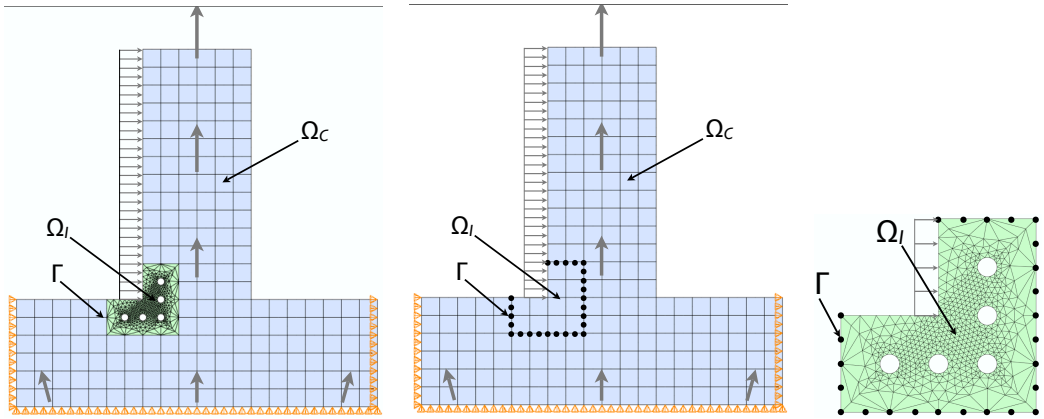

(a) Reference (b) Global (c) Local

Figure 1: Reference, global and local problems

For that purpose, the "global-local" solution $s = (\underline{u}, \underline{\underline{\sigma}})$ (where $\underline{u}$ is the displacement field and $\underline{\underline{\sigma}}$ is the Cauchy stress field) is defined as:

$$s(\underline{x}) = \begin{cases} s^L(\underline{x}) & \text{if } \underline{x} \in \Omega_I \\ s^G(\underline{x}) & \text{if } \underline{x} \in \Omega_C \end{cases} \tag{1}$$

where Ω_I denotes the *area of interest* (described by the local model) and $\Omega_C = \Omega \setminus \Omega_I$ denotes the rest of the structure and is called the *complement area*. Superscripts G and L respectively denote quantities from the global model (defined on Ω in its entirety) or the local model (defined on Ω_I alone). The interface $\partial\Omega_I \cap \mathring{\Omega}$ is denoted by Γ. It is assumed that the local and global meshes are matching on the interface; non-conforming interfaces have not been considered so far, though a standard mortar technique [1] should provide an effective answer to this problem.

It was proved in [9] that the "global-local" solution verifies every equation of the reference problem if, and only if, the following three conditions are satisfied:

(i) the local solution s^L verifies each equation written in Ω_I and on $\partial\Omega_I \setminus \Gamma$ (equilibrium, boundary conditions and elastic-plastic constitutive equations);

(ii) the *restriction* of the global solution to Ω_C, s^G_C, verifies each equation written in Ω_C and on $\partial\Omega_C \setminus \Gamma$ (equilibrium, boundary conditions and linear elastic constitutive equations);

(iii) both solutions match on Γ (displacements are equal and tractions are balanced).

Those three conditions form the so-called *global/local formulation*. They are similar to the formulation of non-overlapping domain decomposition methods: *the overlap of models is eliminated from the equations*.

2.2 Interface formulation

A surface coupling algorithm is used to match the two solutions on Γ: the finite element solvers exchange the nodal interface displacements $\mathbf{u}_\Gamma^G$ and $\mathbf{u}_\Gamma^L$ and their dual quantities, the nodal interface reaction forces $\lambda_{\Gamma,C}^G$ and λ_Γ^L. More specifically, $\lambda_{\Gamma,C}^G$ is the nodal interface force exerted on the *complement area* (hence the subscript C) of the global model through Γ, which is a global internal force, and λ_Γ^L is the nodal interface force exerted on the local model, which is a local external force (since Γ is part of Ω_I's boundary).

It is convenient to write the equations in each subdomain in condensed form, *ie.* using only interface quantities. The global/local formulation can then be translated directly. The global solution $(\mathbf{u}_\Gamma^G, \lambda_{\Gamma,C}^G)$ has to verify equations of Ω_C, that is:

$$\mathbf{S}_{\Gamma,C}^G \mathbf{u}_\Gamma^G - \mathbf{b}_{\Gamma,C}^G = \lambda_{\Gamma,C}^G \tag{2}$$

where $\mathbf{S}_{\Gamma,C}^G$ denotes the Schur complement of the complement domain's stiffness matrix, and $\mathbf{b}_{\Gamma,C}^G$ the corresponding condensed right hand side vector. Likewise, the local solution $(\mathbf{u}_\Gamma^L, \lambda_\Gamma^L)$ has to verify equations of Ω_I, that is:

$$\mathbf{h}_\Gamma^L(\mathbf{u}_\Gamma^L) = \lambda_\Gamma^L \tag{3}$$

where $\mathbf{h}_\Gamma^L$ is the local problem's discrete Steklov-Poincaré operator, *ie.* the interface stiffness operator giving the reaction forces λ_Γ^L as a function of the interface displacements $\mathbf{u}_\Gamma^L$. A different notation is used because this group of equations is nonlinear. Finally, interface displacements must match and interface nodal reaction forces must be balanced:

$$\mathbf{u}_\Gamma^G = \mathbf{u}_\Gamma^L \tag{4}$$

$$\lambda_{\Gamma,C}^G + \lambda_\Gamma^L = 0 \tag{5}$$

This formulation is very similar to nonlinear domain decomposition methods [22]. The difference is that *the global model is never substructured* since the approach is non-intrusive; therefore, the complement area is never separated from the area of interest in the global model, and the quantities $\mathbf{S}_{\Gamma,C}^G$ and $\mathbf{b}_{\Gamma,C}^G$ cannot be accessed directly. Equivalently, we can say that it is impossible to prescribe tractions directly on the complement area through Γ; only *traction discontinuities* between Ω_C and Ω_I can be prescribed in the global model.

The (primal) condensed form of the reference problem would be:

$$\mathbf{h}_\Gamma^L(\mathbf{u}_\Gamma^R) + \mathbf{S}_{\Gamma,C}^G \mathbf{u}_\Gamma^R - \mathbf{b}_{\Gamma,C}^G = 0 \tag{6}$$

It is easy to see that the condensed global/local formulation (2)–(5) implies Equation (6) — that is, that the solution to (2)–(5) is indeed the reference solution.

2.3 Iterative coupling algorithm

The basic algorithm is a modified Newton method formulated on an interface condition: starting from a global displacement field $\mathbf{u}_\Gamma^G$, the following three steps are performed at each iteration. Note that the global forces $\lambda_{\Gamma,C}^G$ are computed, whenever required, according to (2).

1. **Local analysis:** a local solution $(\mathbf{u}_\Gamma^L, \lambda_\Gamma^L)$ is sought, verifying (3) and a boundary condition on Γ; this condition can be prescribed displacements (4), prescribed forces (5) or a combination of both, as it will be detailed in Section 3.

2. **Residual computation:** an interface load vector called the *residual* is computed. It measures the non-verification of interface equations; for example, if prescribed displacements are used, then the residual measures the lack of equilibrium of nodal forces, *ie.* $\mathbf{r}_\Gamma = -(\lambda_{\Gamma,C}^G + \lambda_\Gamma^L)$. Convergence is tested here: if the residual's norm is small enough, iterations are stopped.

3. **Global correction:** the residual is injected into the global problem as an additional interface load. This is done by first solving a corrective global problem loaded only with the residual (all other loads

and boundary conditions being set to zero); in interface notation, introducing $\mathbf{S}_\Gamma^G = \mathbf{S}_{\Gamma,C}^G + \mathbf{S}_{\Gamma,I}^G$ the global model's Schur complement, this writes:

$$\mathbf{S}_\Gamma^G \Delta\mathbf{u}_\Gamma^G = \mathbf{r}_\Gamma \tag{7}$$

Then the global displacement is updated before going back to step 1:

$$\mathbf{u}_\Gamma^G \leftarrow \mathbf{u}_\Gamma^G + \Delta\mathbf{u}_\Gamma^G \tag{8}$$

In [9], prescribed displacements were used as local boundary condition at Step 1; this paper focuses on *mixed* boundary conditions (see Section 3). Note that:

- No Schur complement is computed explicitly: to multiply a Schur complement by a displacement vector means to prescribe this displacement on the interface of the corresponding model, to run the finite element analysis with every other load set to zero, and to collect the corresponding interface reaction forces.
- The basic method is indeed a modified Newton method, since the global stiffness operator is never modified: the global mesh, constitutive relation and the location of Dirichlet boundary conditions are retained (prescribed displacements are set to zero but this does not modify the stiffness matrix) and only additional loads are prescribed. This allows a very simple implementation (see Section 4), since only displacements and forces at the interface nodes are exchanged. This idea was first suggested in some early global/local literature [23]. A drawback is that convergence can be slow in presence of moderate nonlinearity, such as plasticity.
- Convergence acceleration techniques can be introduced. The modified Newton algorithm described above can easily be turned into a quasi-Newton algorithm — namely, the SR1 algorithm (for Symmetric Rank One) — in a completely non-intrusive way. In [9] this technique greatly improved convergence rates for negligible additional costs when using displacement boundary conditions. Section 4 shows that this technique is also efficient with mixed boundary conditions.

3 MIXED BOUNDARY CONDITIONS

The results presented in [9] were obtained using only prescribed displacement conditions. Such conditions are easy to implement, and the global/local displacement field (defined by (1)) obtained through them is continuous; for that reason, they are very popular and, for example, often used in global/local [3, 20, 23] or multi-level FEA [7, 12], as well as hierarchical techniques [5, 8, 21].

However, displacement conditions suffer from a couple of limitations. First, the difference between the two stiffnesses of the area of interest (from the global and local models) often has a great influence on the displacement field, which causes the boundary condition to be inaccurate. In other terms, during the first iterations, global interface displacements are usually quite different from reference interface displacements. This affects the solution's accuracy and, therefore, the algorithm's convergence; a common workaround is to position the interface far enough from the local details, but this increases the size of the local problem unnecessarily. Another problem is that if plasticity occurs at integration points located close to the boundary, we have noticed that the prescribed displacements can cause convergence difficulties.

Mixed boundary conditions are commonly used in domain decomposition methods (see for example [18, 14]) for these reasons. They are known to be quite insensitive to the difference of stiffnesses and give accurate results from the first iterations, which implies fast convergence, and they do not have a tendency to cause stress concentrations or convergence problems. Their increased accuracy was also shown in the field of global/local analysis [13]. Therefore, they seemed to be a natural choice for the method.

3.1 Basic theory

3.1.1 Formulation

The term "mixed boundary condition" can have several meanings. In this paper, we have chosen to use *Robin boundary conditions*. In other words, when we solve the local problem (3), we do not prescribe the

continuity of displacements (4) or the equilibrium of forces (5), but a *linear combination* of those two equations. That is,

$$\mathbf{A}_\Gamma \left(\mathbf{u}_\Gamma^L - \mathbf{u}_\Gamma^G\right) + \left(\lambda_{\Gamma,C}^G + \lambda_\Gamma^L\right) = 0 \tag{9}$$

where $\mathbf{A}_\Gamma$ is a square matrix representing an interface stiffness, and a parameter of the method. Likewise, the residual is not computed on Equations (4) or (5), but also on a linear combination of both:

$$\mathbf{B}_\Gamma \left(\mathbf{u}_\Gamma^L - \mathbf{u}_\Gamma^G\right) - \left(\lambda_{\Gamma,C}^G + \lambda_\Gamma^L\right) = 0 \tag{10}$$

where $\mathbf{B}_\Gamma$ is also a parameter of the method.

It should be clear that Equations (9)–(10) are equivalent to Equations (4)–(5) if, and only if, the matrix $\mathbf{A}_\Gamma + \mathbf{B}_\Gamma$ is invertible. If this condition is verified and if the algorithm is convergent, then the solution will converge to the reference solution, as noted in Section 2.3. Apart from this condition, there is no restriction to the choice of $\mathbf{A}_\Gamma$ and $\mathbf{B}_\Gamma$.

3.1.2 Choosing the interface stiffnesses

However, this choice has a huge impact on the convergence rate — especially the choice of the first stiffness $\mathbf{A}_\Gamma$. This can be seen by combining the local problem's equation (3) with the Robin condition (9); the *mixed local problem* that is obtained is the following:

$$\mathbf{h}_\Gamma^L(\mathbf{u}_\Gamma^L) + \mathbf{A}_\Gamma \mathbf{u}_\Gamma^L = \mathbf{A}_\Gamma \mathbf{u}_\Gamma^G - \lambda_{\Gamma,C}^G \tag{11}$$

In other terms, the interface stiffness $\mathbf{A}_\Gamma$ is *assembled* into the local problem and a mixed load vector $\mathbf{A}_\Gamma \mathbf{u}_\Gamma^G - \lambda_{\Gamma,C}^G$ is prescribed. This shows that the Robin condition can be a powerful tool for *modelling the stiffness* of the rest of the structure, that is the complement domain, as "seen from outside". More specifically, it is easy to see that if $\mathbf{A}_\Gamma$ is equal to the Schur complement of the complement domain $\mathbf{S}_{\Gamma,C}^G$, then the mixed local problem (11) becomes equivalent to the reference problem (6). Therefore, if we could compute $\mathbf{S}_{\Gamma,C}^G$ exactly, the analysis strategy would be:

- *exact* (the local solution would be the exact reference solution),
- *direct* (this would be true at the very first local analysis, so only one iteration would be needed)
- and *completely insensitive to the interface position*, unlike the variant with prescribed displacements.

Unfortunately, to compute $\mathbf{S}_{\Gamma,C}^G$ means to perform static condensation on the complement area, and this technique is known to be extremely expensive on large 3D problems as explained in the introduction (in addition to being intrusive...). Therefore, the Schur complement has to be *approximated*; its approximation should of course be inexpensive to compute, and at the same time be "close enough" to the exact value so that fast convergence is obtained. In addition, since the method is aimed to be non-intrusive, the building of the approximation should work on irregular meshes and require as little user intervention (such as manual mesh operations) as possible.

Many authors have suggested such approximations; in the context of FEA with large subdomains and possibly complex geometries, two common approaches are:

- *short-scale approximations*, usually based on an element strip at the vicinity of the interface;
- *long-scale approximations*, usually inspired from homogenization, *ie.* submitting the subdomain to a small number of "large-scale" interface loads.

In the next subsections, we review those two approaches. We show that when taken individually, none of them brings significant gains in terms of convergence rates. For that reason, we suggest combining them to form a *two-scale approximation*, and we show that such an approximation can lead to very fast convergence.

3.2 A toy problem

To assess those approximations, we consider the toy problem shown on Figure 1, in the previous section. It consists in a simple 2D model with a coarse regular mesh and a linear elastic constitutive law; normal displacements are blocked on the bottom and lower sides edges, and a pressure and a centrifugal loading

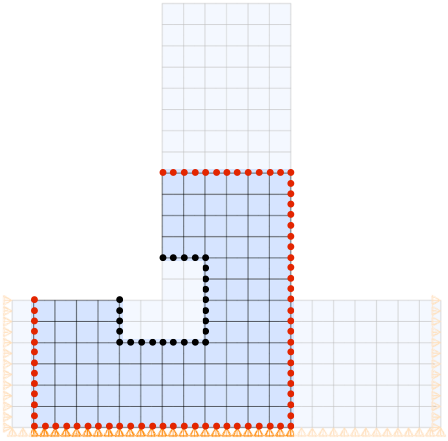

Figure 2: A strip of 4 elements of the complement area adjacent to the interface.

are applied, as shown on Figure 1(b). Since the elastic limit is slightly exceeded near the sharp corner on the left, an elastic-plastic local model is introduced, see Figure 1(c); this model also contains mesh refinements and five small holes that were not present in the global model. The elastic moduli and loading are the same as those of the global model (although they could have been different if needed). Both models are analyzed using Abaqus/Standard, version 6.7.

The quality of the approximations is assessed using two simple indicators, giving a measure of the difference between the local solution and the reference solution. The first indicator measures the global error in interface displacements, *ie.*

$$\eta_{\mathbf{u}} = \frac{||\mathbf{u}_\Gamma^L - \mathbf{u}_\Gamma^R||}{||\mathbf{u}_\Gamma^R||}$$

This indicator can help estimating convergence, since interface displacements are part of the mixed boundary condition. However, they are rarely an interesting quantity for the engineer, and this indicator can be biased since local displacements could very well be close to the reference, whereas local stresses would not; for those reasons, we also use a second indicator based on the maximum equivalent plastic strain:

$$\eta_p = \frac{p_{\max}^L - p_{\max}^R}{p_{\max}^R}$$

This indicator is purely local; in theory, it does not give a reliable measure of the error in the solution, since it could be very low while the local and reference solutions are very different. However, in the case of single-step loads, we have found it to be reliable, and much more sensitive to "coupling errors" (that is, inaccurate boundary conditions) than the displacement indicator. In addition, in applications like turbine blades of aircraft engines, plastic strains are often a quantity of interest (so would stresses, but plastic strains are usually more sensitive, as will be shown in Section 4).

3.3 Element strip techniques

We start by reviewing short-scale, or element strip techniques. As the name suggests, those techniques consist in performing exact static condensation, not on the whole subdomain, but on a strip of elements adjacent to the interface, usually clamped at the other end [18]. This leads to much lower memory requirements. In this study, we have built such "strips" by analyzing element connectivity in a recursive manner (*eg.* a 2-element strip is defined as "elements that touch elements that touch the interface"). This process is simple, scriptable and suitable for irregular meshes. Figure 2 shows an example of a 4-element strip in the complement domain, for the 2D toy problem described above.

Short-scale approaches also include the popular *lumped* and *superlumped* approximations, commonly used as preconditioners in domain decomposition methods. The lumped approximation consists in replacing the whole Schur complement of a matrix with the corresponding interface block; that is, if the

Table 1: Error indicators at the first iteration for "short-scale" approximations

Approximation type	$\eta_{\mathbf{u}}$	η_p
Exact	*0*	*0*
4-element strip	0,122	-0,177
2-element strip	0,133	-0,211
1-element strip	0,140	-0,247
Lumped	0,147	-0,274
Superlumped	0,148	-0,282
Prescribed displacements	*0,153*	*-0,313*

complement area's stiffness matrix (or contribution to the global stiffness matrix) is

$$\mathbf{K}_C^G = \begin{bmatrix} \mathbf{K}_{CC}^G & \mathbf{K}_{C\Gamma}^G \\ \mathbf{K}_{\Gamma C}^G & \mathbf{K}_{\Gamma\Gamma,C}^G \end{bmatrix}$$

then its Schur complement is by definition $\mathbf{S}_{\Gamma,C}^G = \mathbf{K}_{\Gamma\Gamma,C}^G - \mathbf{K}_{\Gamma C}^G \left[\mathbf{K}_{CC}^G\right]^{-1} \mathbf{K}_{C\Gamma}^G$; the corresponding lumped approximation is simply $\mathbf{K}_{\Gamma\Gamma,C}^G$, and the superlumped approximation is its diagonal. Both can be obtained at no additional cost.

To assess those approximations, a single mixed local analysis was performed, starting from the initial global solution; then the two indicators introduced in the previous section were computed. Results are summarized in Table 1. Significant errors can be noticed since, with the best approximation (the 4-element strip of Figure 2), the displacement error is about 12% (versus 15% with prescribed displacements) and the maximum plastic strain error is about 18% (versus 31% with prescribed displacements); the mixed boundary conditions obtained by these approximations do not bring significant gains.

Since the 4-element strip is actually quite large and covers a significant part of the complement domain (see Figure 2), this result suggests that most of the complement domain must be included in the "strip" in order to obtain a realistic local solution. In other terms, the information contained in short-scale approaches may not be sufficient to build efficient approximations at low costs for our problem (although it works well for designing inexpensive preconditioners).

3.4 Homogenization-like techniques

For that reason, we also investigated long-scale approximations, most of which are inspired from multiscale homogenization techniques; that is, to prescribe a small number of interface loads on the whole subdomain, representing their "large-scale" behaviour [12, 14, 10, 11, 7]. In comparison, short-scale techniques consist in prescribing every possible interface load (*ie.* performing condensation) on a small part of the subdomain.

3.4.1 Formulation

Two ways are possible: prescribing displacements and collecting reaction forces will give an approximation to the Schur complement (or primal Schur complement), whereas prescribing forces and collecting displacements will give an approximation to its inverse (or dual Schur complement). We have chosen the first option, for three reasons. First, we are actually interested in the primal Schur complement rather than its inverse. Second, prescribing interface forces on the complement domain would require dissociating it from the area of interest — that is "cutting" the global model in two, since otherwise the internal interface forces would not be controlled (remember, the global model is *not substructured* in our coupling strategy). This is an intrusive operation; on the other hand, prescribing interface displacements does not require such a manipulation ($\lambda_{\Gamma,C}^G$ is then an internal reaction force). Finally, this choice avoids any potential "floating subdomain" problem, that could occur with prescribed forces and require a specific algorithm [19, 6].

In this framework, the "homogenized", approximate Schur complement is defined as:

$$\tilde{\mathbf{S}}_{\Gamma,C}^G = \mathbf{\Pi}\mathbf{S}_{\Gamma,C}^G\mathbf{\Pi}^{\mathrm{T}} \tag{12}$$

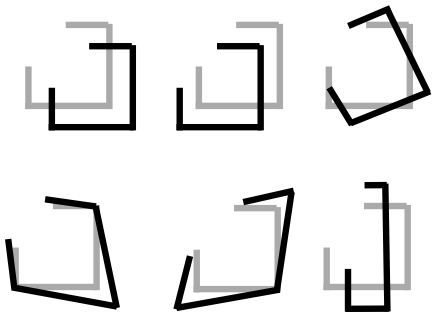

Figure 3: Affine displacement fields on the interface.

where $\mathbf{\Pi}$ is a linear operator *projecting* interface forces onto some low-dimensional "macro" space, and its transpose $\mathbf{\Pi}^{\mathrm{T}}$ does the same with interface displacements. This choice guarantees that the approximation is symmetric. Of course, the matrices in the equation above are never computed explicitly, and the result is obtained by submitting the global model to simultaneous load vectors.

3.4.2 Choosing the "macro" spaces

The efficiency of the method obviously depends on the choice of those "macro" spaces. For "macro" displacements, a common choice is the space of interface rigid body motions (translations and rotations); this allows a convenient definition of "macro" forces by duality arguments. In this study, we have chosen the space of *affine fields* as the "macro" displacement space; this includes translations and rotations, but also stretching and distorting modes, as shown on Figure 3. For the sake of simplicity, we always handle the interface as a whole: decomposing it would probably require user intervention in the case of very irregular meshes.

Concerning "macro" forces, two definitions are possible. The simplest choice, used for example in the LaTIn micro/macro method [14], is to use the *same* "macro" space as for the displacements. If $\mathbf{E}$ is a rectangular matrix whose columns form an orthonormal basis of that space (*ie.* so that $\mathbf{E}^{\mathrm{T}}\mathbf{E} = \mathbf{I}$), the projector can then be defined as:

$$\mathbf{\Pi} = \mathbf{\Pi}^{\mathrm{T}} = \mathbf{E}\mathbf{E}^{\mathrm{T}} \tag{13}$$

that is, as an orthogonal projector (represented as a symmetric matrix).

However, that choice may seem too restrictive because a "macro" space that makes sense for nodal displacements does not necessarily make sense for nodal forces[1]. For that reason, we also considered different "macro" space and projectors for the two quantities (we require, however, their dimensions to be the same). If $\mathbf{E}$ and $\mathbf{F}$ are the rectangular matrices whose columns respectively define the displacement and force "macro" bases, and if we assume those bases are biorthonormal (*ie.* $\mathbf{E}^{\mathrm{T}}\mathbf{F} = \mathbf{F}^{\mathrm{T}}\mathbf{E} = \mathbf{I}$), then the projectors can be defined as:

$$\mathbf{\Pi} = \mathbf{F}\mathbf{E}^{\mathrm{T}} \tag{14}$$

$$\mathbf{\Pi}^{\mathrm{T}} = \mathbf{E}\mathbf{F}^{\mathrm{T}} \tag{15}$$

Since "macro" displacements include unit translations and rotations, the corresponding "macro" forces will have unit resultants and moments. Apart from this, there is no restriction to the choice of the "macro" force space. For reasons that will be explained later, we have chosen to define this space as *the space of reaction forces to prescribed "macro" displacements*. Thus, once these reaction forces have been computed, its basis $\mathbf{F}$ is sought as:

$$\mathbf{F} = \mathbf{S}^{G}_{\Gamma,C}\mathbf{E}\mathbf{P} \tag{16}$$

where $\mathbf{P}$ is a basis transformation matrix; the biorthogonality condition then gives

$$\mathbf{E}^{\mathrm{T}}\mathbf{S}^{G}_{\Gamma,C}\mathbf{E}\mathbf{P} = \mathbf{I} \tag{17}$$

which means $\mathbf{P}$ has to be the inverse or pseudoinverse of $\mathbf{E}^{\mathrm{T}}\mathbf{S}^{G}_{\Gamma,C}\mathbf{E}$. Once this pseudoinverse is computed, the "macro" force basis $\mathbf{F}$ is directly formed by (16). Those "macro" basis vectors do not correspond

[1] For example, a translation motion is represented by a uniform nodal displacement vector, but a uniform traction field, that could be a possible "macro force" candidate, is usually *not* represented by a uniform nodal force vector.

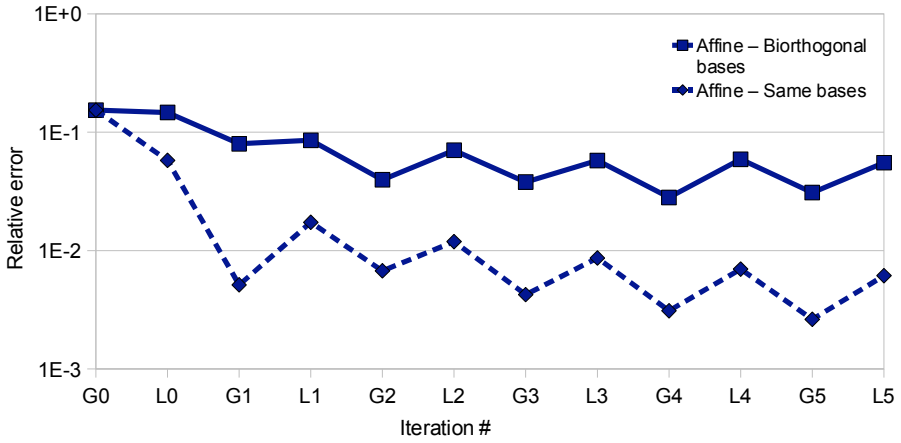

Figure 4: Lack of convergence of the "long-scale" mixed approximation.

to particularly noticeable force distributions; however, they have the desirable properties of having unit resultants or moments (or higher order moments, due to the stretching and distortion terms), and of containing relevant mechanical information on the subdomain's response.

3.4.3 Results

Those two variants were tested on the 2D toy problem described in the previous sections. Figure 4 represents the evolution of the displacement error indicator versus the iteration numbers, for the "basic" algorithm (the letters L and G refer to local and global steps of each iteration); it can be noticed that the indicator stagnates, and does not converge to zero. A more careful examination shows that the error actually *increases* during the local analysis. This shows that the information obtained using this method is not sufficient to create efficient boundary conditions.

3.5 A two-scale Schur complement approximation

The previous subsections have shown that neither short-scale nor long-scale approximations contain enough information to give a good representation of the complement area's influence. It seems necessary to represent *both* the vicinity of the interface and the global behaviour of the subdomain correctly. In this section, we propose an original technique for *combining* those approximations, in order to build an efficient representation of a subdomain's stiffness; it was only applied to the complement domain in the case of our strategy, but could certainly have more applications, for example to building preconditioners in domain decomposition methods.

3.5.1 A consequence of Saint-Venant's principle

This technique is based on *Saint-Venant's principle*, which is a well-known empirical principle commonly used in mechanical engineering in general, and in the beam theory in particular; a discussion is given in [15] for linear elasticity. This principle states that "*... the strains that can be produced in a body by the application, to a small part of its surface, of a system of forces statically equivalent to zero force and zero couple, are of negligible magnitude at distances which are large compared with the linear dimensions of that part*" [17].

In other words, at large distances from the system of forces, the motion of the "body" is almost rigid (since strains are negligible). If we could identify this rigid body motion and substract it from the displacement field, then the residual (non-rigid) displacements would be almost zero. Therefore, clamping that "body" at a sufficient distance from the system of forces should not introduce a significant error in any component of the solution, see Figure 5. This should be true, in particular, for displacements of points where the forces are applied.

Let us apply this principle to the complement domain, subjected to interface forces λ_Γ. Assume that $\mathbf{E}^T\lambda_\Gamma = 0$; since $\mathbf{E}$ contains interface translations and rotations, then λ_Γ's resultant force and moment are zero. Therefore, the principle above states that the complement domain could be replaced with an

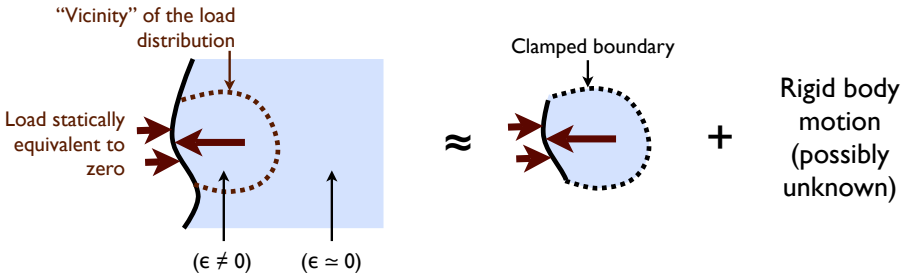

Figure 5: A consequence of Saint-Venant's principle.

element strip such as those presented in Section 3.3, clamped at a sufficient distance from the interface; denoting its Schur complement by $\mathbf{D}_\Gamma$, one should then have:

$$\left[\mathbf{S}^G_{\Gamma,C}\right]^{-1}\lambda_\Gamma \approx \mathbf{D}_\Gamma^{-1}\lambda_\Gamma + \mathbf{E}\mathbf{x}$$

where $\mathbf{x}$ is a (possibly unknown) amplitude vector, and $\mathbf{E}\mathbf{x}$ represents the rigid body motion that was omitted; $\mathbf{x}$ might be zero if each of the complement area's rigid body modes are blocked. Moreover, since every equation written in the complement area is linear, this rigid body motion has to be a linear function of the loading; that is, there exists a rectangular matrix $\mathbf{C}$ such that:

$$\mathbf{x} = \mathbf{C}\lambda_\Gamma$$

Determining this matrix could be difficult in practice, but luckily it will not be needed in the final result. By combining the two equations above, the following postulate is obtained: for any interface load vector λ_Γ,

$$\boxed{\text{if } \mathbf{E}^{\mathrm{T}}\lambda_\Gamma = 0, \text{ then } {\mathbf{S}^G_{C,\Gamma}}^{-1}\lambda_\Gamma \approx \left[\mathbf{D}_\Gamma^{-1} + \mathbf{E}\mathbf{C}\right]\lambda_\Gamma} \tag{18}$$

3.5.2 Interface force splitting

We can now introduce our two-scale approximation to the Schur complement. It basically consists in *splitting, using the projectors defined in Section 3.4.2, the space of interface forces in two:*

- *a "micro" space*, design to contain only loads with zero resultant force and moment, thus that can be analyzed on a short-scale approximation, like those in Section 3.3;
- and *a "macro" space* with a low dimension, that will be analyzed on the whole complement area with homogenization-like techniques, like in Section 3.4.

The formulation is a straightforward application of this idea. First, notice that both definitions of the "macro" force projector $\mathbf{\Pi}$, that are (13) and (14), imply the following property: for any interface force λ_Γ,

$$\mathbf{E}^{\mathrm{T}}\left(\mathbf{I} - \mathbf{\Pi}\right)\lambda_\Gamma = 0 \tag{19}$$

This means that $\left(\mathbf{I} - \mathbf{\Pi}\right)\lambda_\Gamma$*'s resultant force and moment are always zero.* Therefore, if we take an arbitrary λ_Γ and write the following "macro/micro" decomposition:

$$\lambda_\Gamma = \mathbf{\Pi}\lambda_\Gamma + \left(\mathbf{I} - \mathbf{\Pi}\right)\lambda_\Gamma$$

and left-multiply it by $\left[\mathbf{S}^G_{\Gamma,C}\right]^{-1}$, then the postulate (18) gives:

$$\left[\mathbf{S}^G_{\Gamma,C}\right]^{-1}\lambda_\Gamma \approx \left[\mathbf{S}^G_{\Gamma,C}\right]^{-1}\mathbf{\Pi}\lambda_\Gamma + \left[\mathbf{D}_\Gamma^{-1} + \mathbf{E}\mathbf{C}\right]\left(\mathbf{I} - \mathbf{\Pi}\right)\lambda_\Gamma$$

And since this is true for any arbitrary λ, the following approximation is obtained:

$$\left[\mathbf{S}^G_{\Gamma,C}\right]^{-1} \approx \left[\mathbf{S}^G_{\Gamma,C}\right]^{-1}\mathbf{\Pi} + \left[\mathbf{D}_\Gamma^{-1} + \mathbf{E}\mathbf{C}\right]\left(\mathbf{I} - \mathbf{\Pi}\right) \tag{20}$$

At first glance, this definition seems to involve the response of the complement domain to prescribed "macro" forces, due to the $\left[\mathbf{S}^G_{\Gamma,C}\right]^{-1}\mathbf{\Pi}$ term; as explained in the previous section, prescribing interface

forces on the complement domain is an intrusive operation. However, this difficulty can be lifted by choosing the "macro" force space accordingly — that is, by defining it as the space of reaction forces to prescribed "macro" displacements, as specified by Equations (14)–(17). With this definition, once those reactions forces have been computed, obtaining $\left[\mathbf{S}^G_{\Gamma,C}\right]^{-1}\mathbf{\Pi}$ is no longer a problem. In addition, Equations (14), (16) and (17) give the following identities:

$$\begin{aligned}\left[\mathbf{S}^G_{\Gamma,C}\right]^{-1}\mathbf{\Pi} &= \mathbf{\Pi}^{\mathrm{T}}\left[\mathbf{S}^G_{\Gamma,C}\right]^{-1} \\ &= \mathbf{\Pi}^{\mathrm{T}}\left[\mathbf{S}^G_{\Gamma,C}\right]^{-1}\mathbf{\Pi} \\ &= \mathbf{E}\left[\mathbf{E}^{\mathrm{T}}\mathbf{S}^G_{\Gamma,C}\mathbf{E}\right]^{-1}\mathbf{E}^{\mathrm{T}}\end{aligned} \tag{21}$$

that is, the response to prescribed "macro" forces is always a "macro" displacement, and the response to a prescribed "macro" displacement is always "macro" forces. In other terms, the following decomposition holds:

$$\left[\mathbf{S}^G_{\Gamma,C}\right]^{-1} \;=\; \mathbf{\Pi}^{\mathrm{T}}\left[\mathbf{S}^G_{\Gamma,C}\right]^{-1}\mathbf{\Pi} \;+\; (\mathbf{I}-\mathbf{\Pi})^{\mathrm{T}}\left[\mathbf{S}^G_{\Gamma,C}\right]^{-1}(\mathbf{I}-\mathbf{\Pi}) \tag{22}$$

This decomposition is a direct consequence of the relation between the "macro" force and displacement spaces, and does not hold with other definitions (*ie.* when the two spaces are chosen to be the same).

Another difficulty is that the expression in (20) looks *unsymmetric*, and thus seems to violate a fundamental property of linear elasticity. More precisely, we *know* that the actual Schur complement is symmetric in our problem, therefore it is desirable to have a symmetric approximation. However, the identities (21) show that the "macro" term is already symmetric, and only the "micro" term (with $(\mathbf{I}-\mathbf{\Pi})$) is unsymmetric. Several ways of making it symmetric can be thought of; similarly to the decomposition (22), we have chosen to slightly adapt Equation (20) as follows:

$$\left[\mathbf{S}^G_{\Gamma,C}\right]^{-1} \;\approx\; \mathbf{\Pi}^{\mathrm{T}}\left[\mathbf{S}^G_{\Gamma,C}\right]^{-1}\mathbf{\Pi} \;+\; (\mathbf{I}-\mathbf{\Pi})^{\mathrm{T}}\left[\mathbf{D}^{-1}_{\Gamma}+\mathbf{E}\mathbf{C}\right](\mathbf{I}-\mathbf{\Pi}) \tag{23}$$

Only the "micro" term changes between the two equations; this change consists in using the short-scale approximation to calculate the "micro" response to "micro" forces only. Note that this expression is still a valid consequence of Saint-Venant's principle: when comparing it to (22), which is an *exact* equation, only the response to "micro" forces is approximated.

Finally, taking the transpose of Equation (19) shows that the rigid body term (with $\mathbf{EC}$) vanishes. Thus there is no need to identify rigid body modes and compute the matrix $\mathbf{C}$, and the final, symmetric *two-scale Schur complement approximation* is:

$$\boxed{\left[\mathbf{S}^G_{\Gamma,C}\right]^{-1} \;\approx\; \mathbf{\Pi}^{\mathrm{T}}\left[\mathbf{S}^G_{\Gamma,C}\right]^{-1}\mathbf{\Pi} \;+\; (\mathbf{I}-\mathbf{\Pi})^{\mathrm{T}}\mathbf{D}^{-1}_{\Gamma}(\mathbf{I}-\mathbf{\Pi})} \tag{24}$$

In practice, the inverse of this equation is computed by writing it as $\mathbf{D}^{-1}_{\Gamma}$ plus a small-rank corrective term, and using Woodbury's matrix identity to obtain directly an approximation to the primal Schur complement. The total cost of the operation is quite low, since it requires:

- a global analysis with 6 simultaneous load vectors (for the "long-scale" part), and
- full condensation of a small element-strip model (for the "short-scale" part)

which is usually much cheaper than performing full condensation of the complement domain, at least on large 3D problems.

3.5.3 Results

The values of the two error indicators obtained after the first local analysis, on the 2D toy problem, are given in Table 2. The "long-scale" part uses prescribed affine displacements and biorthonormal bases, as explained previously; for the "short-scale" part, the same approximations as those of Section 3.3 were tested. Local results appear to be much better than with the long-scale or short-scale approximations alone, and errors are 3 to 6 times lower than with the displacement condition. This proves that combining information obtained from both scales is a promising way for building realistic approximations to the Schur complement (or stiffness) of a subdomain.

In Figure 6, the corresponding local solution (Figure 6(a)) is compared to the local solution obtained with prescribed displacements (Figure 6(c)), again at the first iteration and in terms of equivalent plastic

Table 2: Error indicators at the first iteration for two-scale approximations

Short-scale term	$\eta_{\mathbf{u}}$	η_p
Exact	*0*	*0*
4-element strip	0,028	-0,050
2-element strip	0,032	-0,054
1-element strip	0,039	-0,064
Lumped	0,047	-0,076
Superlumped	0,051	-0,085
Using prescribed displacements	*0,153*	*-0,313*

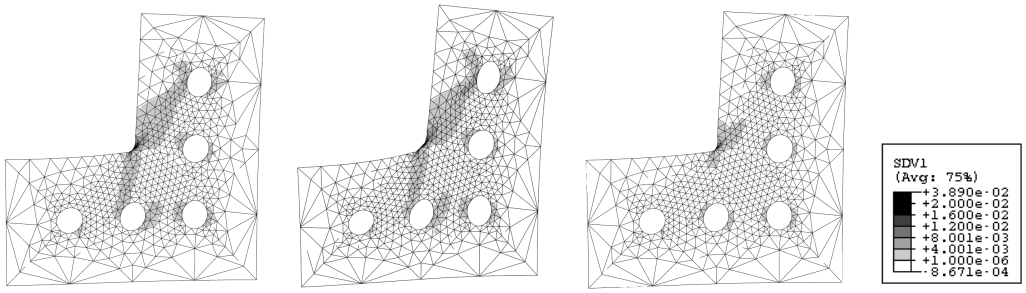

(a) First local iteration, two-scale mixed condition

(b) Reference

(c) First local iteration, displacement condition

Figure 6: Contour plots of equivalent plastic strains obtained with various methods

strains: it is clearly much closer to the reference solution (Figure 6(b)). Therefore, mixed boundary conditions that use this technique can lead to improved local solutions, at realistic costs.

On a side note, this two-scale approximation can certainly be used in other applications than the strategy described above; it can probably be useful in the context of preconditioning or mixed domain decomposition methods, in particular when the subdomains are large.

3.6 Convergence results

Convergence of the plastic strain error indicator is represented on Figure 7, for the four variants of the method. Grey curves correspond to the prescribed displacement conditions (also called *primal*, as in domain decomposition methods), and black curves to mixed conditions; dotted curves were obtained using the "basic" (*ie.* modified Newton) algorithm, and plain curves using the "accelerated" (*ie.* quasi-Newton) algorithm. Iterations are numbered so that the first global and local analyses, which correspond to a weak "one-way" coupling, correspond to iteration zero; thus they give the number of *additional analyses* carried out by the strategy.

The mixed boundary condition seems to bring significant improvements in convergence rates, especially with the basic algorithm; an error of 10^{-3} on the maximum plastic strain (which is a quite strict threshold, by engineering standards) is reached after 2 iterations using the mixed condition, versus 7 iterations using the primal condition. Gains in terms of iteration numbers are less evident with the accelerated algorithm (2 iterations versus 3), but it still can be noticed that the error indicator decreases faster. Last, convergence acceleration does not impact the mixed variant's speed as much as it does to the primal variant (see also [9]), but it still provides a significant increase in convergence rates for negligible additional costs.

4 AN INDUSTRIAL EXAMPLE

Finally, the feasibility and performance of this variant were assessed by running additional tests on a larger 3D test case provided by Snecma. This case consists of two Abaqus/Standard models represented on Figure 8 and is based on an aircraft engine's turbine blade, with an extremely simplified geometry for

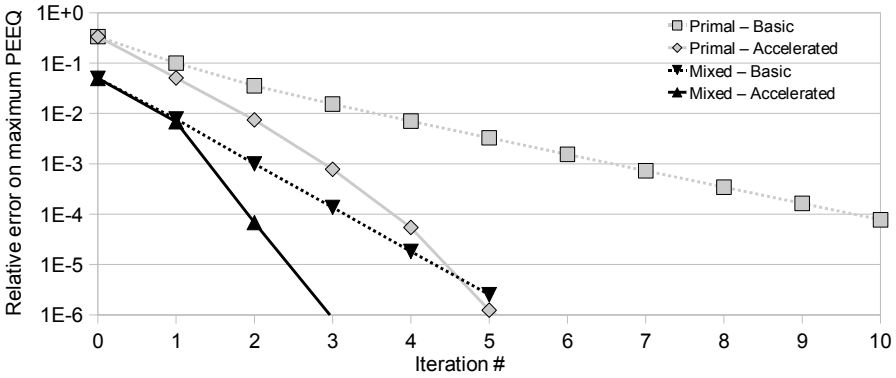

Figure 7: Convergence of maximum plastic strains on the 2D example

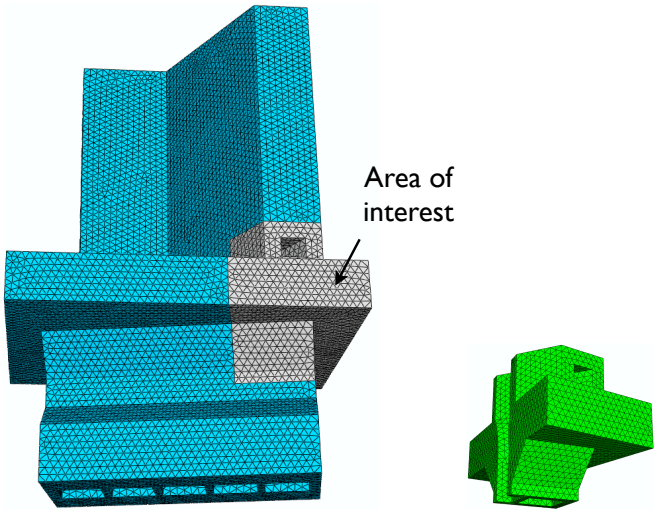

Figure 8: Global and local meshes of the 3D test case

confidentiality purposes. The global model is subjected to a complex set of loads that includes centrifugal forces, a pressure field on the blade's surface, and a precomputed temperature field; normal displacements are blocked on the slanted surfaces at the bottom, and the constitutive law is linear elastic. It contains about 500,000 DOFs. The local model is located at the bottom of the blade and contains the plastic area, with reasonable margins; it features an elastic-plastic constitutive law and its mesh is locally identical to the global mesh (no refinements or geometric details were introduced). It contains 80,000 DOFs. The interface consists of three planes and totals 6,400 DOFs.

4.1 Implementation

The two models had initially been designed to be analyzed with *submodeling*, which is a very common global/local analysis procedure available in most major FEA software, and is basically equivalent to the first iteration of our method (as a consequence, it only ensures an incomplete, "one-way" coupling, since the global influence of the local modifications are ignored). Since the strategy is designed to be modular and non-intrusive as explained in Section 2, only a few lines had to be modified in the input files in order to adapt them to the strategy.

The algorithm itself was coded as a simple *Python script* of a few hundred lines, reading results from output databases using Abaqus' built-in methods, processing them, writing boundary conditions or additional loads to text files, and submitting jobs. For the mixed boundary condition, scripts were also used to build the "macro" problems (from the node coordinate table) and the element strip models (from the element connectivity table), to analyze them and to compute the approximate Schur complement automatically, with almost no user intervention. Typically, starting from an operational submodeling data set

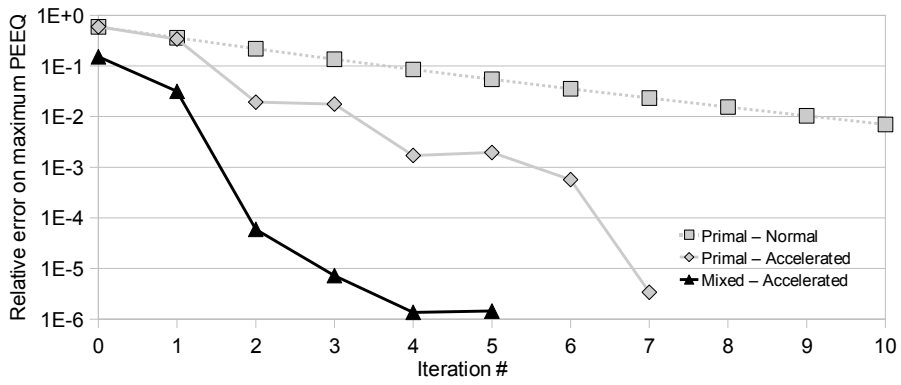

Figure 9: Convergence of maximum plastic strains on the 3D example

and tested Python scripts, preparing the models for the strategy could take a few minutes (displacement variant) to one hour (mixed variant), which seems quite affordable; this is the main advantage of using a non-intrusive framework.

4.2 Submodeling errors and convergence

Results are shown on Figure 9; three variants of the method were tested. Notice the high initial value of the indicator for the two "primal" curves, of about 60%; this initial value corresponds to the first local solution, that would be obtained by traditional submodeling. This means that ignoring the global influence of local plasticity leads to underestimating the maximum equivalent plastic strain by 60%, which is probably unacceptable if this quantity is important for the application.

Of course, this test case is quite different from actual turbine blade problems since it uses a monotonic load and an elastic-plastic constitutive law, whereas turbine blades are subjected to cyclic loads and exhibit important viscous effects. However, this example shows that submodeling can actually introduce huge errors on some important output quantities, even when more common quantities such as stresses do not seem to be affected (the relative error on the maximum Von Mises stress is much smaller, about 0.5%). Once again, the mixed boundary condition is quite efficient in reducing this error; at "iteration zero", an error of about 13% is obtained, which is probably not sufficient but is 5 times better than with the displacement condition.

When the method is carried out until convergence, the basic primal algorithm appears to be quite slow, as expected: it takes 9 iterations to bring the plastic strain error indicator down to 10^{-2}, which is a common threshold for engineering standards. The accelerated primal algorithm takes 4 iterations before reaching that level, whereas the additional cost per iteration is completely negligible. The "mixed" results were obtained using the two-scale Schur complement approximation; computing this approximation consists in:

- performing static condensation on a 3 element wide strip ("short-scale" part)
- running a global computation with 12 simultaneous right-hand sides, corresponding to the prescribed affine "macro" displacements ("long-scale" part)

Then the algorithm takes 2 iterations only to reach the threshold, which shows once again the efficiency of the two-scale approximation.

However notice that the proposed approximation to the Schur complement is a dense matrix whose handling is CPU-time consuming, and that Abaqus/Standard (6.7-1) comes with limitations on the reuse of the global stiffness matrix which penalize our strategy. As a consequence, on this single increment case, our strategy uses more CPU time than the direct solution of the reference problem. We are currently investigating solutions that should reduce the cost per iteration while maintaining high convergence rates.

5 CONCLUSIONS

In this paper, a two-scale Schur complement approximation was introduced. By combining short-scale and long-scale information, this approximation can give a realistic representation of the stiffness of a large subdomain at realistic costs. It was derived and tested in the framework of a non-intrusive coupling strategy, that consisted in introducing local plasticity and geometric details in a small region of a large finite element problem by the means of an additional local model; in this paper, this local model used mixed boundary conditions, and using a two-scale approximation of the Schur complement of the rest of the structure led to very fast convergence, whereas short-scale or long-scale approximations were not sufficient by themselves. This approximation can probably be useful in various situations found in domain decomposition methods, such as preconditioning or building mixed interface conditions.

Concerning the coupling strategy, it has currently been tested on problems limited to one single load increment. To handle multiple increments, two possibilities are foreseen: either incrementation is performed and exchanges are carried out at each increment (like in traditional implicit coupling schemes) or, conversely, exchanges are performed and each global or local problem contains a whole loading history. The latter option, called *non-incremental*, seems more promising for several reasons. First, it would be more convenient to implement around FEA software such as Abaqus, which is designed to analyze complete loading histories on a given problem, and not just one specific increment: one would only have to exchange data between whole analysis jobs, whereas implicit incremental coupling schemes are not supported in the version we use (Abaqus/Standard 6.7-1). Second, the non-incremental approach minimizes information transfers between the solvers, and reduces the number of decompositions of the global stiffness matrix, which are time-consuming (during the global correction step, the right-hand side vectors corresponding to the residuals at every time step could be processed *at the same time*). Finally, it would easily allow using different time steps for the two models. In order to do this, the formulation and algorithm should be adapted to a non-incremental form, such as those used in the LaTIn method (see for example [2]).

Apart from this, we envision three extensions of the method. One is about handling *non-matching meshes* between the two models, with an appropriate technique such as the mortar method [1]; this would certainly allow the strategy to address a more general class of locally nonlinear problems, and to be a more flexible structural reanalysis and model coupling tool. Another is about handling *two different solvers*; namely, a commercial FEA package for the global problem, coupled with in-house software for the local problem. This would allow testing advanced models or solution techniques, that are almost impossible to implement in commercial software at the moment, on complex industrial problems that cannot be solved with "research software" alone; a possible example would be the inclusion of localized damage on a large aircraft part made out of laminated composites. Finally, many more configurations can be thought of, such as multiple areas of interest (*ie.* several local models) or multiple zooming levels (*ie.* more than two); this could be useful in a number of engineering situations and give more reliability to the usual "global-to-local" simulation scheme, used in many industrial applications of computational mechanics.

Acknowledgment: This work has been supported by Snecma.

References

[1] Bernardi C, Maday Y, Patera AT. A new nonconforming approach to domain decomposition: the mortar element method. In *Nonlinear partial differential equations and their applications, Collège de France seminar*, Brézis H, Lions JL (eds), Pitman: London, 1991; 13–51,

[2] Boucard P.A, Ladevèze P, Poss M, Rougée P. A nonincremental approach for large displacement problems. *Computers and Structures* 1997; **64(1–4)**:499–508 .

[3] Cormier NG, Smallwood BS, Sinclair GB, Meda G. Aggressive submodelling of stress concentrations. *International Journal for Numerical Methods in Engineering* 1999; **46**:889–909.

[4] Cresta P, Allix O, Rey C, Guinard S. Nonlinear localization strategies for domain decomposition methods: application to post-buckling analyses. *Computer Methods in Applied Mechanics and Engineering* 2007; **196(8)**:1436–1446.

[5] Düster A, Rank E, Steinl G, Wunderlich W. A combination of an h- and a p-version of the finite element method for elastic-plastic problems. In *ECCM '99. CD-ROM proceedings of the European Conference on Computational Mechanics*, Wunderlich W (ed). Munich, 1999.

[6] Farhat C, Roux FX. Implicit parallel processing in structural mechanics. *Computational Mechanics Advances* 1994; **2(1)**:1-124.

[7] Feyel F and Chaboche J.L. FE^2 multiscale approach for modelling the elastoviscoplastic behaviour of long fiber SiC/Ti composite materials. *Computer Methods in Applied Mechanics and Engineering* 2000; **126**:17–38.

[8] Fish J. The *s*-version of the finite element method. *Computers and Structures* 1992; **43(3)**:539–547.

[9] Gendre L, Allix O, Gosselet P, Comte F. Non-intrusive and exact global/local techniques for structural problems with local plasticity. *Computational Mechanics* 2009; **44**:233–245.

[10] Guidault PA, Allix O, Champaney L, Navarro JP. A two-scale approach with homogenization for the computation of cracked structures. *Computers and Structures* 2007; **85(17–18)**:1360-1371.

[11] Guidault PA, Allix O, Champaney L, Cornuault C. A multiscale extended finite element method for crack propagation. *Computer Methods in Applied Mechanics and Engineering* 2007; **197(5)**:381–399.

[12] Ibrahimbegovic A, Markovic D. Strong coupling methods in multi-phase and multi-scale modeling of inelastic behavior of heterogeneous structures. *Computer Methods in Applied Mechanics and Engineering* 2003; **192**:3089–3107.

[13] Jara-Almonte CC, Knight CE. The specified boundary stiffness and force (SBSF) method for finite element subregion analysis. *International Journal for Numerical Methods in Engineering* 1988; **26**:1567–1578.

[14] Ladevèze P, Dureisseix D. A micro/macro approach for parallel computing of heterogeneous structures. *International Journal for Computational Civil and Structural Engineering* 2000; **1**:18–28.

[15] Ladevèze P, Simmonds J. New concepts for linear beam theory with arbitrary geometry and loading. *European Journal of Mechanics - A/Solids* 1998; **17(3)**:377–402.

[16] Le Tallec P. Domain decomposition methods in computational mechanics. *Computational Mechanics Advances* 1994; **1(2)**:121-220.

[17] Love AEH, *A treatise on the mathematical theory of elasticity*. Cambridge University Press, 1927. (Dover reprint ISBN 0486601749)

[18] Magoulès F, Roux FX, Series L. Algebraic approximations of Dirichlet-to-Neumann maps for the equations of linear elasticity. *Computer Methods in Applied Mechanics and Engineering* 2006; **195**:3742–3759.

[19] Mandel J, Dohrmann CR. Convergence of a balancing domain decomposition by constraints and energy minimization, *Numer. Linear Algebra Appl.* 2003; **10**:639–659.

[20] Mao KM, Sun CT. A refined global-local finite element analysis method. *International Journal for Numerical Methods in Engineering* 1991; **32**:29–43.

[21] Oden JT, Zohdi TI. Analysis and adaptive modeling of highly heterogeneous structures. *Computer Methods in Applied Mechanics and Engineering* 1997; **148**:367–392.

[22] Pebrel J, Rey C, Gosselet P. A nonlinear dual domain decomposition method: application to structural problems with damage. *International Journal of Multiscale Computational Engineering* 2008; **6**:251–262.

[23] Whitcomb JD. Iterative global-local finite element analysis, *Computers and Structures* 1991; **40(4)**:1027–1031.