Coupled Systems Mechanics, *Vol. 8, No. 2 (2019) 99–109*
DOI: https://doi.org/10.12989/csm.2019.8.2.099

A FE^2 multi-scale implementation for modeling composite materials on distributed architectures

Guido Giuntoli*[1], Jimmy Aguilar[1], Mariano Vázquez[1], Sergio Oller[2] and Guillaume Houzeaux[1]

[1]*Computer Application in Science & Engineering Department, Barcelona Supercomputing Center, Carrer de Jordi Girona, 29, 31, 08034 Barcelona, Spain*
[2]*International Centre for Numerical Methods in Engineering, Universitat Politècnica de Catalunya, Carrer del Gran Capità, S/N, 08034 Barcelona, Spain*

(Received September 30, 2018, Revised February 24, 2019, Accepted February 25, 2019)

Abstract. This work investigates the accuracy and performance of a FE^2 multi-scale implementation used to predict the behavior of composite materials. The equations are formulated assuming the small deformations solid mechanics approach in non-linear material models with hardening plasticity. The uniform strain boundary conditions are applied for the macro-to-micro transitions. A parallel algorithm was implemented in order to solve large engineering problems. The scheme proposed takes advantage of the domain decomposition method at the macro-scale and the coupling between each subdomain with a micro-scale model. The precision of the method is validated with a composite material problem and scalability tests are performed for showing the efficiency.

Keywords: FE^2; multi-scale; HPC; composite materials

1. Introduction

One of the major challenges in today's composite material design is in the prediction of the strength and reliability of the structures. Composites are widely used nowadays for building components in different industries. The A350XWB, a commercial aircraft of Airbus, is composed with 53% by weight in composite materials (see Mrazova 2013). The main reason is that composites offer admirable mechanical properties such as a high strength-weight relation and excellent fatigue, corrosion and impact resistance. Moreover, the reduction in weight also decreases the fuel consumption and the CO_2 emissions. The use of composites materials is expected to increase and, for making an efficient use of them, the optimization of the structures will be necessary in the design stage. For achieving this, accurate and efficient numerical algorithms are needed to solve this problems in a reasonable time.

The FE^2 multi-scale method was developed with the purpose of solving large heterogeneous material problems that present a strong difference between the length scales (see Suquet 1987),

*Corresponding author, E-mail: guido.giuntoli@bsc.es

http://www.techno-press.org/?journal=csm&subpage=8 ISSN: 2234-2184(Print), 2234-2192(Online)

i.e., the radius of a carbon fiber and the length of a composite material panel. Generally real size problems are not feasible to be solved by applying directly the finite element (FE) method because the computational cost is extremely large. Other techniques such as the phenomenological models like the mixture theory developed by Truesdell (1960) are generally computationally affordable but are not easy to be developed for complex microstructures. The FE^2 multi-scale method is designed to avoid the use of phenomenological model by solving two coupled FE problems: one in the macro-scale and the other in the micro-scale. The method reduces enormously the computational cost of the single scale FE method but it continues to be computational expensive for real size problems.

To tackle this bottleneck, an efficient FE^2 multi-scale algorithm was implemented with a fully parallelized philosophy. The method was developed assuming small deformations on the solid mechanics equations and making use of the uniform strain boundary condition for the macro-to-micro transitions. The computational implementation is based on a parallel algorithm that applies the domain decomposition method at the macro-scale level and couples each subdomain with an independent micro-scale code that calculates the average properties for a Gauss point group. Non-linear composite material problems with hardening plasticity are studied in order to measure the accuracy of the FE^2 method. The solutions are compared with the single scale FE method to quantify the boundary effects on a periodic structure. On the other hand, strong scalability tests are performed in order to show the efficiency of the parallel strategy implemented. Different mesh sizes for the macro-scale and the micro-scale are studied.

The paper is organized as follows. In Sec. 2 the FE^2 method and the governing equations are presented. In Sec. 3 the parallel computational strategy is explained. Then, Sec. 4 discusses the computational results related with the accuracy and performance. Finally, in Sec. 5 the main conclusions are exposed.

2. FE^2 algorithm

The FE^2 multi-scale method assumes that the original heterogeneous problem can be solved by a homogeneous FE model at the macro-scale level and a microscopic FE model capable of predicting the average properties of a representative volume element (RVE) for being used at the macro-scale (see Suquet 1987). The main unknowns and variables of interest in solid mechanics problems are the stress field tensor σ, the displacement field u and the deformation field tensor ϵ. Knowing the value of them it is possible to predict if a composite material structure is feasible to be built.

At the macro-scale level the governing equations used are the solid mechanics equilibrium ones undergoing the small deformation approach. The partial derivative equation (PDE) subjected to boundary conditions and coupled with the constitutive laws reads

$$\begin{cases} \nabla \cdot \overline{\sigma} = 0 & \text{in } \overline{\Omega}, \\ \overline{u} = \overline{u}_d & \text{in } \overline{x} \in \overline{\Gamma}_d, \\ \overline{\sigma} \cdot \hat{n} = \overline{\sigma}_n \cdot \hat{n} & \text{in } \overline{x} \in \overline{\Gamma}_n, \\ \overline{\sigma} = \langle \sigma \rangle. \end{cases} \tag{1}$$

In the last, $\overline{\Omega}$ is the macroscopic domain and $\overline{\Gamma}_d$ and $\overline{\Gamma}_n$ are the boundary sections where Dirichlet and Neumann condition are applied respectively. The constitutive material law ($\overline{\sigma} = \langle \sigma \rangle$) is not strictly defined, the stress $\overline{\sigma}$ is the average of the microscopic stress field σ. This last, is a result of

the *localization* of the macroscopic strain $\overline{\epsilon}$ into the micro-scale problem.

For the micro-scale level problem, it should be assumed that the macroscopic variables are an average of the microscopic ones. Considering that Ω represents the domain of a RVE at the micro-scale level, the macroscopic strain and stress are calculated as

$$\overline{\epsilon} = \frac{1}{\text{vol}(\Omega)} \int_{\Omega} \epsilon \, dV \quad \text{and} \quad \overline{\sigma} = \frac{1}{\text{vol}(\Omega)} \int_{\Omega} \sigma \, dV. \tag{2}$$

The governing equation's assumptions for the micro-scale are the same as in the macro-scale. The PDE, joined with the constitutive material laws and the boundary conditions define the system

$$\begin{cases} \nabla \cdot \sigma = 0 & \text{in } \Omega, \\ u = \overline{\epsilon} \cdot x & \text{in } x \in \Gamma, \\ \sigma = f(\epsilon, q). \end{cases} \tag{3}$$

In this case the constitutive laws for all the materials present in the microstructure is defined by $\sigma = f(\epsilon, q)$ where q are the internal variables used to model the non-linear materials. The boundary condition $u = \overline{\epsilon} \cdot x$ is called *uniform strain*, and makes Eq. (2) to be satisfied. Other types, such as the *uniform stress* or the *periodic* boundary conditions also satisfy Eq. (2), for that they are called *admissible* boundary conditions (see Miehe 2002).

The process of solving the system represented by Eq. (3), applying the boundary condition that depends on $\overline{\epsilon}$, is called *localization* . On the other hand, the calculation of the macroscopic stress tensor $\overline{\sigma}$ applying the last of Eq. (1) is called *homogenization*. The FE^2 scheme is visualized in Fig. 1.

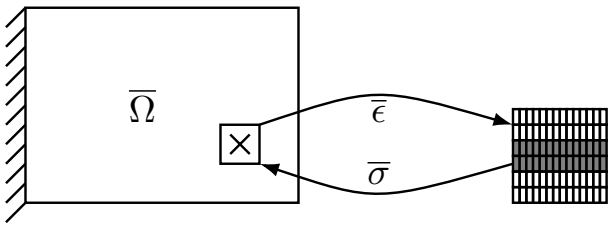

Fig. 1 Scheme of the FE^2 multi-scale method

The FE^2 method considers to solve the problems represented by Eqs. (1) and (3) both by the FE method. A brief explanation is given now following the work of Johnson (1960). The procedure is the same for the macro-scale and the micro-scale problems so the notation that follows is generic.

First, the strong form of the differential equations $\nabla \cdot \sigma = 0$ is transformed into a weak form that reads:

$$r(u) = \int_{\Omega} \sigma(u) : \delta\epsilon \, dV = 0. \tag{4}$$

The Neumann (natural) boundary condition term is omitted here for simplicity. This expression is called *virtual work principle* and $\delta\epsilon$ is called *virtual strain*.

The FE method assumes that the solution can be approximated as an interpolation of a finite number of values and base functions that are defined for a FE mesh used to approximate the domain. This is defined as

$$\tilde{u}(x) = \sum_{i}^{N} u_i \psi_i(x), \tag{5}$$

The index i relates the node in the mesh and N is the total number of them. On the other hand, $\psi_i(x)$ are the shape functions used to interpolate the displacement field defined at the nodes (u_i). Then, substituting Eq. (5) into Eq. (4) a finite system of equations is obtained

$$r_j(\tilde{u}) = \int_{\Omega} \sigma(\tilde{u}) : \delta\epsilon_j \; dV = 0 \quad \text{for} \quad j = 1 \dots N. \tag{6}$$

To solve the last non-linear system of equations, the multi-dimensional Newton-Raphson iterative procedure is used. The algorithm is stated as

$$\begin{aligned} du &= -\left[\tfrac{dr}{du}\right]^{-1} r(u_{n+1}^k), \\ u_{n+1}^{k+1} &= u_{n+1}^k + du. \end{aligned} \tag{7}$$

The method begins with an initial guess $u_{n+1}^0 = u_n$ where u_n is the solution converged in the previous time step. The $(N \times \text{dim}) \times (N \times \text{dim})$ matrix $\mathbb{A} = \frac{dr}{du}$ is called *jacobian* matrix and it is sparse due to the definitions of the shape functions ψ_i. The iterative procedure ends when the norm of the residue ($||r(u_{n+1}^k)||$) is below a small tolerance that depends on each problem.

The process for assembly the jacobian matrix and the residue is related with an integral in the whole domain. The assembly process can be done going through each of the elements of the mesh and calculating an elemental matrix $\mathbb{A}_e$ and an elemental vector r_e. These values are added in the corresponding positions of the global matrix $\mathbb{A}$ and residue r as

$$\mathbb{A} = \sum_{e}^{n_e} \mathbb{A}_e \;\text{ and }\; r = \sum_{e}^{n_e} r_e. \tag{8}$$

The elemental residue r_e for every element is obtained by the gaussian quadrature rule

$$r_e = \sum_{g}^{n_g} \mathbb{B}^T \sigma \; w_g. \tag{9}$$

The matrix $\mathbb{B}$ is built with the derivatives of the shape functions $\psi(x)$ as it is explained by Bathe (2015). On the other hand, the values w_g correspond to the weights of the integration quadrature. In a similar way, the elemental matrix $\mathbb{A}_e$ is calculated by

$$\mathbb{A}_e = \sum_{g}^{n_g} \mathbb{B}^T \mathbb{C} \; \mathbb{B} \; w_g. \tag{10}$$

Where $\mathbb{C}$ is the tangent constitutive tensor, a property that depends on the material and it is defined as

$$\mathbb{C} = \frac{d\sigma}{d\epsilon}. \tag{11}$$

When the constitutive material law $\sigma = f(\epsilon, q)$ has a simple expression, it is possible to find an analytic expression for $\mathbb{C}$. However, in most of the cases and in the macro-scale level problem this is not possible to be done, therefore, a numerical perturbation procedure can be used to estimate this value. It consists in applying small perturbations $\delta\epsilon$ over the strain ϵ for each of its six components $i = 1 \dots 6$ (three components in two dimension problems)

$$\epsilon_j^* = \begin{cases} \epsilon_i & \text{if } i \neq j, \\ \epsilon_i + \delta\epsilon & \text{if } i = j. \end{cases} \tag{12}$$

For the micro-scale problem, the stress tensor σ^* can be evaluated using the constitutive law $\sigma^* = f(\epsilon^*)$; and for the macroscopic case, the stress $\overline{\sigma}^*$ can be homogenized by localizing $\overline{\epsilon}^*$, as it was explained. By this way, each column j of $\mathbb{C}$ can be estimated by

$$\mathbb{C}_j \approx \frac{\sigma^* - \sigma}{\delta\epsilon}, \tag{13}$$

that corresponds to the two point derivative approximation.

3. Computational implementation

As it was explained, the FE problem in any of the scales has two dominant processes. The first is the assembly of $\mathbb{A}$ and r for creating the jacobian matrix and the residue, this algorithm is $\mathcal{O}(N)$ where N is the number of elements in the mesh. The second is to solve the linear system of equations which has a computational cost between $\mathcal{O}(N)$ and $\mathcal{O}(N^3)$ (see Johnson 1960). This means, that for large problems the solving stage is the dominant process.

In the case of the complete FE^2 multi-scale algorithm the time needed to perform the calculation is the time to assembly and solve the macroscopic problem. Considering that N_e is the number of elements at the macro-scale, $\mathcal{O}(S_\mu)$ and $\mathcal{O}(S_M)$ are the costs to solve a microscopic and macroscopic problem, respectively; the cost of the algorithm can be approximated as

$$\mathcal{O}(\text{FE}^2) \approx \mathcal{O}(N_e\ \mathcal{O}(S_\mu)) + \mathcal{O}(S_M). \tag{14}$$

It is important to remark that when the microscopic problem is large enough, the total time can be dominated by it and to be proportional to the number of elements at the macro-scale: $\mathcal{O}(N_e\ \mathcal{O}(S_\mu))$. In the opposite case, when the problem at the micro-scale level is small enough, the FE^2 algorithm behaves as a single scale FE method: $\mathcal{O}(S_M)$. For both scenarios, a fully parallelizable algorithm should be implemented in order to deal with real size problems and to achieve reasonable computation time.

The solution applied in this work is based on the domain decomposition method at the macro-scale level. Basically, the method decomposes the spatial domain into smaller parts and the total work for the assembly and the solving stages is distributed among a group of processes (see Dolean 2015).

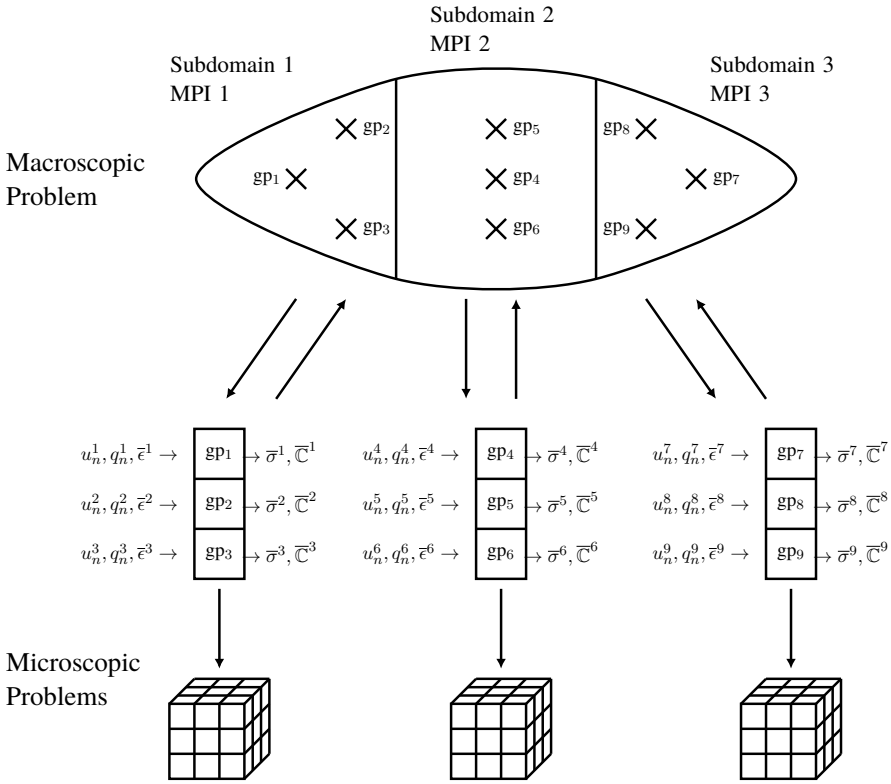

Fig. 2 Computational scheme for the FE^2 multi-scale method. The macroscopic problem applies the domain decomposition technique to solve the FE problem and each subdomain has a microscopic code associated responsible of calculating the properties at each Gauss point of the subdomain

The macro-scale decomposition is feasible because the computation of each microscopic problem is independent from the others and no communication exists between them.

The computational scheme is visualized on Fig. 2. Each macroscopic subdomain is communicated with an independent micro-scale algorithm which is in charge of computing, for each Gauss point, the average properties $\overline{\sigma}$ and $\overline{\mathbb{C}}$ using $\overline{\epsilon}$ and also the fields of the previous time step: the displacement u_n and the internal variables q_n.

The parallel multi-physics FE code Alya (see Vázquez 2014) was used to solve the macro-scale problem. On the other hand, MicroPP, a free sequential C++ library (see MicroPP 2018), was implemented from scratch to solve the microscopic part. The basic algorithm used at the macro-scale is outlined in Alg. 1, also the procedure for the calculation of $\overline{\sigma}$ at the micro-scale is shown in Alg. 2.

Alg 1: FE algorithm for the macro-scale problem.

for $t \leftarrow 1$ **to** N_t **do**
  $\overline{u} = \overline{u}_d(t)$
  while $||\overline{r}|| > tol$ **do**
    for $\overline{e} \leftarrow 1$ **to** N_e^{sub} **do**
      for $\overline{g} \leftarrow 1$ **to** N_g **do**
        micro $\rightarrow$ **SetStrain**(gp$_i$, $\overline{\epsilon}$)
      end
    end
    micro $\rightarrow$ **Homogenize**()
    for $\overline{e} \leftarrow 1$ **to** N_e^{sub} **do**
      for $\overline{g} \leftarrow 1$ **to** N_g **do**
        micro $\rightarrow$ **GetStress**(gp$_i$, $\overline{\sigma}$)
        micro $\rightarrow$ **GetCtan**(gp$_i$, $\overline{\mathbb{C}}$)
        $\overline{r}_e = \overline{r}_e + \overline{\mathbb{B}}^T\ \overline{\sigma}$ and $\overline{\mathbb{A}}_e = \overline{\mathbb{A}}_e + \overline{\mathbb{B}}^T\ \overline{\mathbb{C}}\ \overline{\mathbb{B}}$
      end
      $\overline{r} = \overline{r} + \overline{r}_e$ and $\overline{\mathbb{A}} = \overline{\mathbb{A}} + \overline{\mathbb{A}}_e$
    end
    $d\overline{u} = -\overline{\mathbb{A}}^{-1}\overline{r}$
    $\overline{u}_{n+1} = \overline{u}_n + d\overline{u}$
  end
  micro $\rightarrow$ **Update**()
end

SetStrain : sends $\overline{\epsilon}$ for a Gauss point to the microscopic code.

Homogenize : localizes the strains $\overline{\epsilon}$ for all Gauss points and stores $\overline{\sigma}$ and $\overline{\mathbb{C}}$.

GetStress & GetCtan : returns $\overline{\sigma}$ and $\overline{\mathbb{C}}$ for a Gauss point.

Update : updates the internal variables q_n and the displacements u_n for all the Gauss points.

Alg 2: FE algorithm for calculating $\overline{\sigma}$ at the micro-scale.

$u = \overline{\epsilon} \cdot x$
while $||r|| > tol$ **do**
  for $e \leftarrow 1$ **to** n_e **do**
    for $g \leftarrow 1$ **to** N_g **do**
      $r_e = r_e + \mathbb{B}^t\ \sigma(\epsilon)$ and $\mathbb{A}_e = \mathbb{A}_e + \mathbb{B}^t\ \mathbb{C}(\epsilon)\ \mathbb{B}$
    end
    $r = r + r_e$ and $\mathbb{A} = \mathbb{A} + \mathbb{A}_e$
  end
  $du = -\mathbb{A}^{-1}r$
  $u_{n+1} = u_n + du$
end
$\overline{\sigma} = \frac{1}{V}\int_V \sigma dV$

4. Results and discussion

To compare the results of the FE^2 multi-scale method, three problems that consist in a beam with two, four and eight layers were solved using a one-scale FE method. These are indicated on the left of Fig. 3. On the right of Fig. 3, the problems at both scales using the FE^2 method are outline. These multi-layer problems aim to quantify the boundary effects not consider in the FE^2 method (see Hollister 1992). The layers alternate elastic and plastic materials. The proportion of each material is also varied for: 0%, 33%, 50%, 66% and 100% for plastic composition. The linear isotropic model is used for the elastic material and a Von-Mises yield criteria combined with an isotropic hardening law explained by Simo (1998) is used to model the plastic one. The constants used are: Young module $E = 3.0 \times 10^7\,\mathrm{N\,m^{-2}}$, Poisson module $\nu = 0.25$, yield stress $\sigma_y = 1.0 \times 10^5\,\mathrm{N\,m^{-2}}$ and hardening parameter $K_a = 2.0 \times 10^5\,\mathrm{N\,m^{-2}}$. On the other hand, the beam's dimensions are $L_x = 10\mathrm{m}$, $L_y = 1\mathrm{m}$ and $L_z = 1\mathrm{m}$.

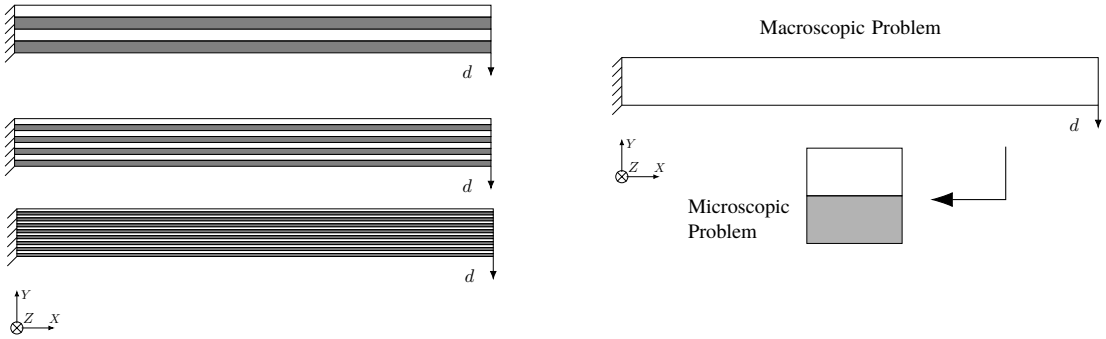

Fig. 3 On the left, three geometries with two, four, and eight layers solved with the single scale FE method. On the right, the problems for the macro-scale and micro-scale solved with the FE^2 method. The white color represents the elastic material and the grey the plastic one

A mesh convergence analysis was done to determine the number of elements needed to get a representative solution. The mesh selected for the three problems with the single scale FE method has $34 \times 16 \times 1 = 544$ elements in the x, y and z directions, respectively. For the FE^2 model, the macro-scale mesh has $30 \times 3 \times 2 = 180$ elements and the microscopic problem has a mesh of $3 \times 3 \times 3 = 27$ elements.

The boundary conditions applied to the beam are

$$\begin{cases} \overline{u} = 0 & & \text{if } x = 0, \\ \overline{u} = \begin{cases} -1.6\,t & \text{if } 0 \le t \le 0.5 \\ -0.8 + 1.6\,(t - 0.5) & \text{if } t < 1.0 \end{cases} & & \text{if } x = L_x, \\ \overline{\sigma} \cdot \hat{n} = 0 & & \text{if } y = 0, y = L_y, z = 0 \text{ or } z = L_z. \end{cases} \tag{15}$$

Finally, the comparison magnitude used is the force in the surface $x = L_x$ defined by

$$F = \int_{x=L_x} \overline{\sigma} \cdot \hat{n}\, dS. \tag{16}$$

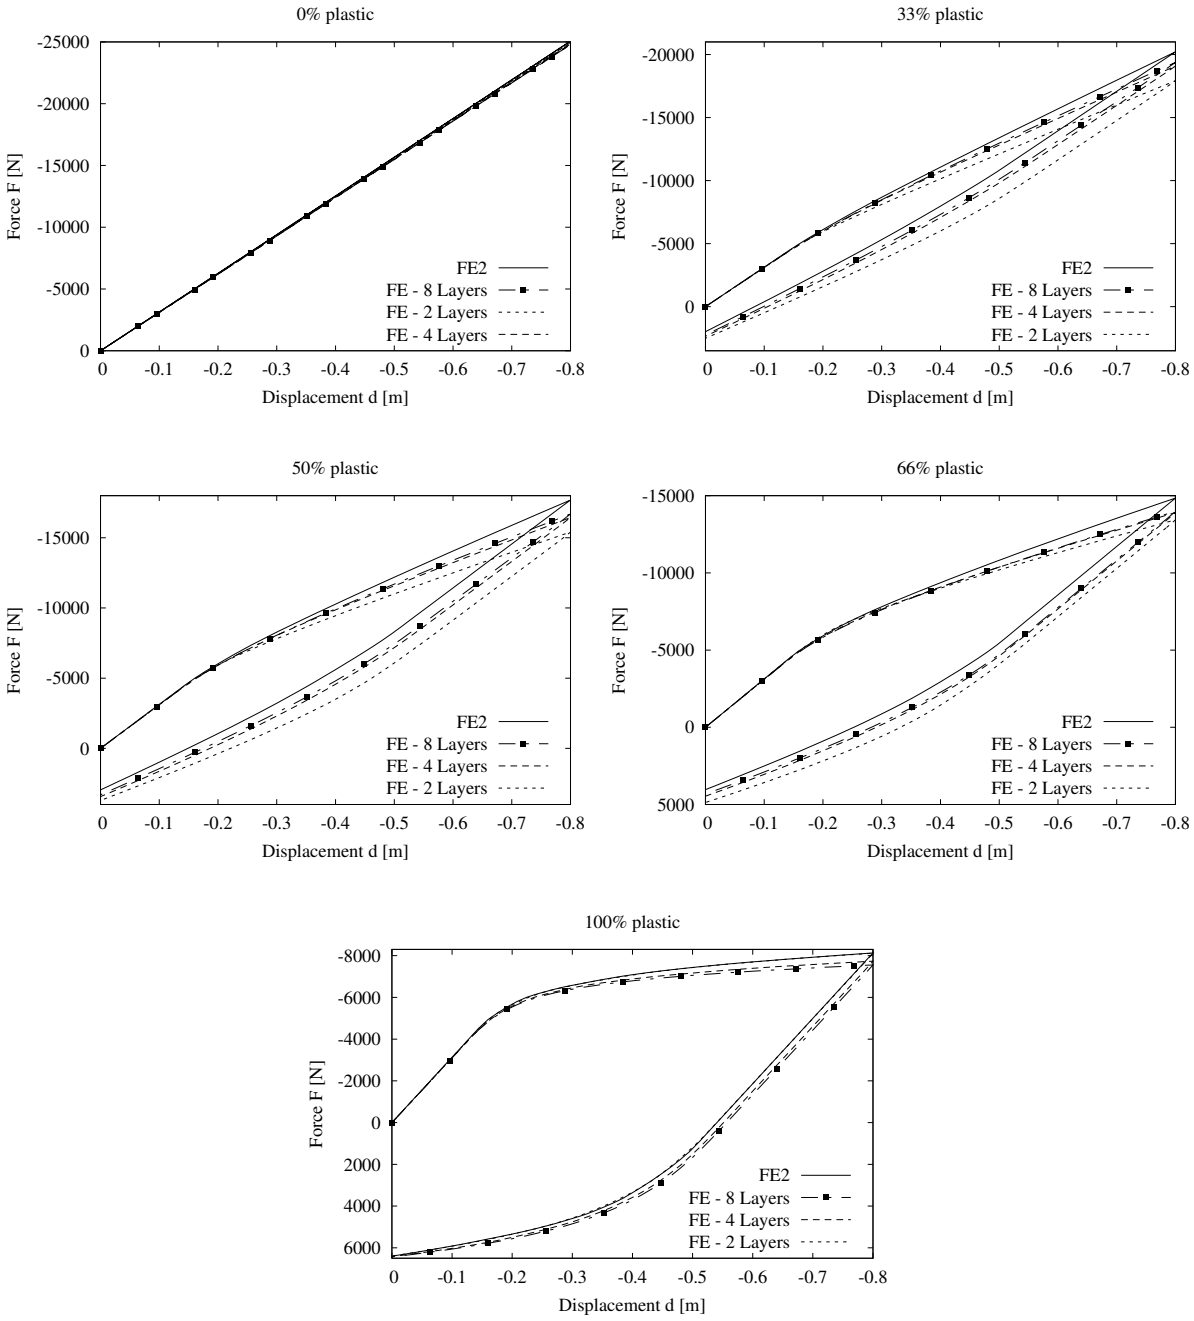

Fig. 4 Results obtained applying the single scale FE and the FE^2 multi-scale method

4.1 Accuracy

The comparisons between the single scale FE method and the FE^2 multi-scale method are visualized in Fig. 4 for different plastic compositions. In all the cases, the solution obtained by the FE^2 method is more similar to the eight layer problem and presents a more remarkable difference in the two layers case. The boundary effect is less significant as the number of layer increases because a larger fraction of the beam volume is actually deformed like the boundary condition model assumed for the micro-scale problem.

4.2 Performance

To test the performance of the algorithm, strong performance tests were executed on Marenostrum IV supercomputer. This architecture consists in 3456 nodes of 2 sockets Intel Xeon Platinum 8160 CPU with 24 cores each working at 2.10 GHz. The cache memory levels are L1 32 Kb, L2 1024 Kb and L3 of $33\,792$ Kb. The compilation of the code was done with GCC/7.1.0 and for the MPI communications the Intel MPI Library 2017 Update 3 was used. The tests consist in five time steps each where the material remained on the elastic zone. By this way, each time steps was solved applying one Newton-Raphson iteration to achieved the equilibrium.

In first place, the computation times varying the micro-scale's mesh were measured. This result is visualized on the left of Fig. 5. The speedup factor increases as the microscopic mesh size grows. This happens because more computational work is assigned to each macroscopic process at the assembly stage, this stage demands less communication time between the MPI processes than the solving stage and this is finally translated in a better performance.

In second place, the results shown on the right of Fig. 5 were performed varying the macroscopic mesh size. As the number of elements increases also the performance grows because each MPI process has more elements and more work is assigned to the assembly and solving stages while the communications increases at a slower rate.

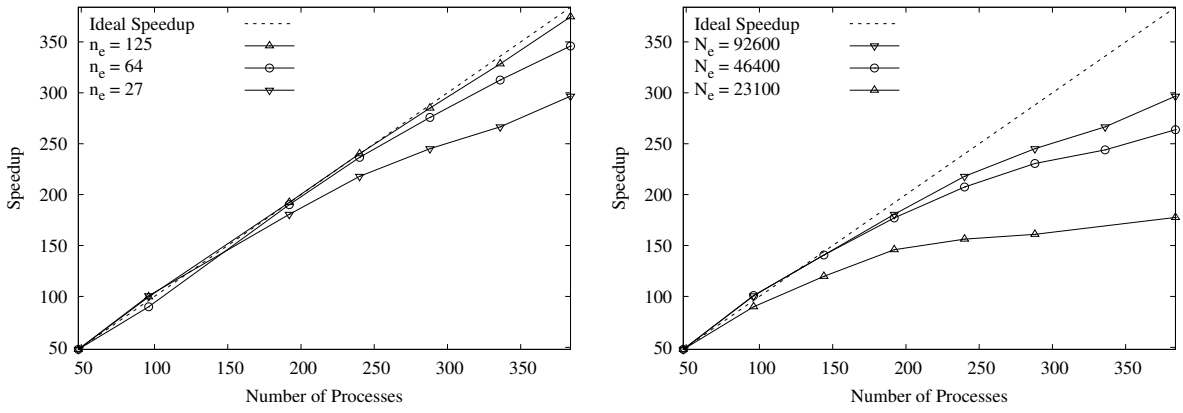

Fig. 5 Strong scalability tests of the FE^2 implementation in Marenostrum IV supercomputer. On the left, microscopic mesh's sizes of 27, 64 and 125 elements are tested using a macroscopic mesh of $92\,600$ elements. On the right, macroscopic mesh's sizes of $92\,600$, $46\,400$ and $23\,100$ elements are tested with a microscopic mesh of 27 elements

5. Conclusions

The FE^2 multi-scale implementation predicted similar results than the one-scale FE method for non-linear materials with hardening plasticity undergoing small deformations. The comparison made for the beam problem, varying the number of material layers and solved with the single scale FE method, showed that the FE^2 multi-scale method is more accurate for periodic media cases where the boundary effects are less significant. It is also important to remark that the experiences, varying the percentage of elastic and plastic compositions from extreme values to intermediate ones, shown that the FE^2 multi-scale method is flexible and capable of modelling a wide spectra of non-linear microstructures.

The parallel computational scheme designed proved to be scalable to large problems. The speedup increases as the meshes in both scales grow because more work is done by each macroscopic process while the communications between them grows in less magnitude.

Additional considerations should be taken for non-linear problems that exhibit a softening response, these include material damage models cases. The convergence of these problems is generally more difficult to achieve and requires more advanced numerical methods. In the work presented by Hautefeuille (2011) a procedure based on the arc-length technique is implemented to calculate the "snap-back" response of those type of systems. This topic is out-of-the-scope of the present article but it is important to mention that the parallel algorithm presented would have a similar speedup that the one shown for the non-linear examples.

References

Aguilar, J. and Giuntoli, G. (2018), *MicroPP: Reference Manual*, <https://github.com/GG1991/micropp>.

Bathe, K.J. (2014), *Finite Element Procedures*, 2nd Edition.

Dolean, V., Jolivet, P. and Nataf, F. (2015), *An Introduction to Domain Decomposition Methods: algorithms, theory and parallel implementation*, SIAM.

Hautefeuille, M., Colliat, J.B., Ibrahimbegovic, A., Matthies, H.G. and Villon, P. (2012), "A multi-scale approach to model localized failure with softening".*Comput. Struct.*, **94**, 83-95.

Hollister, S.J. and Kikuchi, N. (1992), "A comparison of homogenization and standard mechanics analyses for periodic porous composites", *Comput. Mech.*, **10**(2), 73-95.

Johnson, C. (1987), *Numerical Solution of Partial Differential Equations by the Finite Element Method*, Dover.

Miehe, C. and Koch A. (2002), "Computational micro-to-macro transitions of discretized microstructures undergoing small strains", *Arch. Appl. Mech.*, **72**(4-5), 300-317.

Mrazova, M. (2013), *Advanced Composite Materials of the Future in Aerospace Industry*.

Simo, J.C. and Hughes T.J.R. (1998), *Computational Inelasticity*, Springer.

Suquet, P.M. (1987), *Elements of Homogenization for Inelastic Solid Mechanics*, Springer-Verlag.

Truesdell, C. and Toupin, R. (1960), *The Classical Field Theories*, Springer-Verlag.

Vázquez M., Houzeaux, G., Koric, S., Artigues, A., Aguado-Sierra, J., Aris, R., Mira, D., Calmet, H., Cucchietti, F., Owen, H., Taha, A. and Cela, J.M. (2014), *Alya: Towards Exascale for Engineering Simulation Codes*, Cornell University, U.S.A.