11th World Congress on Computational Mechanics (WCCM XI)
5th European Conference on Computational Mechanics (ECCM V)
6th European Conference on Computational Fluid Dynamics (ECFD VI)
E. Oñate, J. Oliver and A. Huerta (Eds)

DOMAIN DECOMPOSITION PRECONDITIONING FOR NON-LINEAR ELASTICITY PROBLEMS

Eirik Keilegavlen*, Jan Ole Skogestad and Jan M. Nordbotten

University of Bergen
Allégaten 41, 5008, Bergen, Norway
Eirik.Keilegavlen@math.uib.no (Presenting author)
Jan.Skogestad@math.uib.no
Jan.Nordbotten@math.uib.no

Key Words: *Non-linear preconditioning, ASPIN, non-linear elasticity, domain decomposition, Newton methods.*

Abstract. We consider domain decomposition techniques for a non-linear elasticity problem. Our main focus is on non-linear preconditioning, realized in the framework of additive Schwarz preconditioned inexact Newton (ASPIN) methods. The standard 1-level ASPIN method is extended to a 2-level method by adding a non-linear coarse solver. Numerical experiments show that the coarse component is necessary for scalability in terms of linear iterations inside the Newton loop. Moreover, for problems that are dominated by non-linearities that are not localized in space the non-linear coarse iterations are crucial for achieving computational efficiency.

1 INTRODUCTION

Physical systems that exhibit non-linear behavior are encountered throughout the sciences. As a result, discretized numerical simulations of practical applications commonly involve solving systems of non-linear algebraic equations. The standard approach is to apply a Newton linearization, and solve the resulting linear system with either a direct or an iterative solver. For large systems of equations the computational cost of the Newton iterations and their associated linear systems of equations can dominate the overall simulation effort and efficient approaches for solving non-linear equations is thus an active research field.

For many problems non-linearities will often be localized in space. This may be caused by the character of the governing equations, heterogeneities in parameter fields, or by boundary conditions. For these problems, a natural approach is to resolve the non-linearities locally instead of by global computations. Domain decomposition (DD) techniques [1] not only allow for such localized treatment of the computational challenges, but also provide a convenient framework for parallel computing, which is necessary to deal with large-scale problems. However, it is well-known that DD is inferior as a stand-alone solver, and in practice the standard Newton method combined with a Krylov subspace linear solver with an appropriate preconditioner, possibly linear DD, is the preferred option. To be competitive, non-linear DD must similarly be considered as a preconditioner, as is realized in the framework of the Additive Schwarz Preconditioned Inexact Newton (ASPIN) methods [2]. ASPIN has been shown to be a robust of efficient approach for solving non-linear problems for many

applications.

To be suited for large-scale parallel computations, the DD method must be scalable when the number of subdomains is increased. To that end, the fine-scale computations need to be amended by a coarse-scale solver, and this can be realized either by deriving non-linear coarse equations, or by instead solving a representative linear coarse problem. The latter option will transfer more of the computational burden to the local problems, but for some problems it may still be an attractive approach in terms of overall computational time.

In this paper, we employ ASPIN to non-linear systems of equations arising from non-linear elasticity problems, with focus on the design of coarse solvers. We show that for this problem, the localization of non-linearities, and thus the optimal design of the coarse spaces, depends critically on the problem setup and parameter field. The paper is structured as follows: The governing equation and its discretization are introduced in Section 2. Section 3 presents solution techniques for non-linear systems. Numerical examples are given in Section 4, and conclusions are drawn in Section 5.

2 GOVERNING EQUATIONS AND DISCRETIZATION

We consider the equilibrium configuration of an elastic medium in $\Omega \in R^d$, with fixed Dirichlet boundary $\partial\Omega_D$ of positive measure and Neumann boundary $\partial\Omega_N = \partial\Omega \setminus \partial\Omega_N$. Introducing H_0^1 as the space of weakly differential functions with zero trace on $\partial\Omega_D$, we can express the solution through the weak formulation [3]: Find $u \in \left(H_0^1(\Omega)\right)^d$ such that

$$(\epsilon(\phi), \sigma(u)) = (g, \phi) \quad \forall \phi \in \left(H_0^1(\Omega)\right)^d \tag{1}$$

Here, the standard L^2 inner product on Ω is denoted $(\,,)$, g represents body forces, while the stress tensor σ is taken to be a non-linear extension of the isotropic stress tensor:

$$\sigma_{i,j}(u) = 2\mu\epsilon(u)_{i,j} + \lambda\, tr(\epsilon(u))\delta_{i,j} + \alpha\epsilon(u)^{\beta} \tag{2}$$

with $\epsilon(u)$ being the symmetrized gradient of the deformation u;

$$\epsilon(u) = \frac{\nabla u + \nabla u^T}{2}$$

We note that the standard linear elasticity problem is recovered by setting $\alpha = 0$.

We next consider a finite element discretization of (1). To that end, Ω is discretized by a triangulation Ω_h, and we denote by V_h the space of nodal basis functions for $P_1(\Omega_h) \subset H_0^1(\Omega)$. The discrete problem then reads: Find $u_h \in (V_h)^d$ such that

$$\left(\epsilon(\phi_h), \sigma(u_h)\right) = (g, \phi_h) \quad \forall \phi_h \in (V_h)^d \tag{3}$$

This is a non-linear problem in terms of the displacement u_h, and the rest of the paper is devoted to solution techniques for this equation.

3 NON-LINEAR PRECONDITIONING TECHNIQUES

We write the non-linear equation (3) in the generic form

$$F(u) = 0 \tag{4}$$

For simplicity we will drop the subscript h throughout this section. The goal of this section is to introduce a two-level ASPIN solver for (4), but we first give a brief introduction of the single level ASPIN method.

3.1 1-level ASPIN methods

Divide the domain into N subdomains Ω_i, which are allowed to overlap. To simplify the construction of the coarse operator (see next section), we will assume that Ω_i can be identified as essentially triangles or quadrilaterals, in the sense that 3 or 4 vertices can be identified in a meaningful way. Let n denote the total number of degrees of freedom in (5), and let n_i be the number of unknowns associated with subdomain i. Restriction operators $R_i \in R^{n_i \times n}$ and prolongation operators $P_i \in R^{n \times n_i}$ are then defined for each subdomain. In this work we let element (j,k) of R_i be unity if the unknown j of subdomain i is associated with the global unknown k, and set $P_i = R_i^T$, but the formulation is valid also for more general approaches [4].

With the above notation, a local non-linear problem can be stated as follows: Starting with the global state u^n, find a local response $\delta_i(u^n)$ such that

$$F_i(\delta_i(u^n); u^n) = R_i F\big(u^n + P_i \delta_i(u^n)\big) = 0, \quad i = 1, .., N \tag{5}$$

Since the boundary conditions for the local problem are inexact, it is computationally efficient to apply an inexact solution strategy to (5). The global update $\delta(u^n)$ is then given as

$$\delta(u^n) = \sum_{i=1}^{N} w_i \delta_i(u^n) \tag{6}$$

where the weight in the vectors w_i handle overlapping subdomains, and are unity in the interior of the subdomains. A domain decomposition iterative solver can now be realized by defining $u^{n+1} = u^n + \delta(u^n)$, where the update is computed from (5)-(6). However a much more efficient approach is found by instead solving the preconditioned non-linear system

$$\delta(u^n) = 0 \tag{7}$$

for u^n using Newton's method [2]. The Jacobian matrix of $\delta(u^n)$, which can be computed from (5)-(6), depend on both u^n and the local updates $\delta_i(u^n)$, and it is thus both cumbersome to form the matrix, and costly to solve the associated linear system. In terms of computational efficiency, it is better to use an approximated Jacobian matrix solve the system [2], [5]

$$\sum_{i=1}^{N} J_i^{-1}(u^n) J(u^n) s^n = \delta(u^n) \tag{8}$$

for the Newton search direction s^n. Here $J(u^n)$ is the Jacobian matrix of $F(u^n)$, i.e. without accounting for the local updates and $\sum_{i=1}^{N} J_i^{-1}(u^n)$ represents the linear additive Schwarz preconditioner. The system is solved by GMRES, with a slight modification to ensure that only the left hand side is preconditioned; the right hand side is already preconditioned non-linearly [2]. Again, we only seek an inexact solution of (8). Finally, as a safeguard in cases where we are far from the solution of (4), the global state is updated as $u^{n+1} = u^n + \lambda^n s^n$,

where the damping parameter λ^n is computed using a line search with cubic backtracking [6].

3.2 Two-level ASPIN

We next turn our attention to the formulation of a coarse system to complement the local systems defined by (5) and the corresponding restriction and prolongation operators. Introduce a coarse grid Ω_0 with the subdomains Ω_i as cells. With the goal of obtaining a coarse-scale discretization that is reminiscent of the fine-scale problem (3), coarse degrees of freedom are assigned with all vertexes of Ω_0. Let n_0 denote the number of coarse unknowns. Construct first-order vertex-centered basis functions ϕ_H, that is linear (on triangles) or bi-linear (on quadrilaterals) functions, on the coarse grid. The restriction operator $R_0 \in R^{n_0 \times n}$ is then defined as a matrix having one function in ϕ_H in each of its rows, and $P_0 = R_0^T$.

With these coarse restriction and prolongation matrices the coarse non-linear system is now defined by setting $i = 0$ in (4). That is, the coarse system of equations is given by

$$F_0(\delta_0(u^n); u^n) = R_0 F\big(u^n + P_0\delta_0(u^n)\big) = 0 \tag{9}$$

and the coarse Jacobian matrix is thus given by

$$J_0(\delta_0(u^n); u^n) = R_0 J\big(u^n + \delta_0(u^n)\big)P_0 \tag{10}$$

Note that with the current formulation of the coarse system, all upscaling or multiscale modeling must be realized in the restriction and prolongation operators.

The coarse solver is introduced to the above 1-level in a multiplicative step, in accordance with [7]. For convenience, we summarize the 2-level ASPIN method in Table 1, and note that the 1-level method is recovered by dropping all references to the coarse level 0.

Table 1: Summary of a single global Newton iteration with the 2-level ASPIN method

1. Solve local non-linear problems
   a. Coarse update: Solve (9) for $\delta_0(u^n)$ with a relative tolerance of τ_{nl}^{rel}, and define a temporary fine-scale state $u^{n,*} = u^n + P_0\delta_0(u^n)$
   b. Local fine-scale updates: Solve the local systems $F_i(\delta_i(u^{n,*}); u^{n,*}) = 0, i = 1,..,N$ for $\delta_i(u^{n,*})$, and define the global update

$$\delta(u^n) = \delta_0(u^n) + \sum_{i=1}^{N} w_i\delta_i(u^{n,*})$$

   Check convergence for $\delta(u^n)$.
2. Within a relative tolerance of τ_{lin}^{rel}, solve the global linear system

$$\left(\sum_{i=1}^{N} J_i^{-1}(u^n)\right) J_0^{-1}(u^n)J(u^n)s^n = \delta(u^n)$$

3. Define the new global state $u^{n+1} = u^n + \lambda^n s^n$, where the damping parameter λ^n is found by cubic backtracking.

4 NUMERICAL RESULTS

In this section we probe the above formulation of a two-level ASPIN method with two numerical test examples. In the first test scalability is tested on a problem with homogenous parameter field. The second case considers a heterogeneous permeability field which is design to explore the effects of localization of non-linearities. Both domains are triangulated, with about 40.000 vertexes. The subdomain geometry will be described in each case. In all cases, an overlap of two vertexes is assigned, and a Restricted Additive Schwarz-type handling of the overlap is applied [8]. All non-linear systems are solved with an exact evaluation of the Jacobian matrix, using the automatic differentiation framework from [9]. The relative tolerances $\boldsymbol{\tau_{nl}^{rel}}$ and $\boldsymbol{\tau_{lin}^{rel}}$ are both set to $\mathbf{10^{-6}}$.

There are two possible benefits from amending the 1-level ASPIN method by a non-linear coarse solver. Firstly, in cases where the non-linearities are not localized, the coarse non-linear solver will contribute to reduce the number of global Newton iterations. Secondly, the coarse solver will contribute not only in the non-linear solves in (4), but also as a coarse component of the linear additive Schwarz preconditioner of (7). Motivated by this, we compare the 2-level ASPIN method by two other solution approaches: A 1-level ASPIN method, as described in section 3.1, and a standard Newton solver, where the inner linear system is solved by GMRES preconditioned by a linear 2-level DD method, i.e. the same preconditioner as is employed for the 2-level ASPIN method. For all methods the setup in terms of the geometry of subdomains and tolerances for iterative solver is identical. Three metrics are employed to measure performance: The number of global Newton iterations, the total number of GMRES iterations, and the number of local linear solves. The latter is composed of the linear preconditioning, and for the ASPIN methods also the solution of local non-linear systems, thus it gives an indication of whether the added cost in solving local non-linear systems is compensated by a correspondingly reduction of linear solver cost. For all problems considered, the coarse systems are small, thus the cost of solving them is negligible.

4.1 A homogeneous non-linear media

In the first test, the computational domain is the unit square. The elastic parameters are homogeneous, with values $\mu = 10^3, \lambda = 10^2, \alpha = 10^4, \beta = 6$, see (2). The displacement field has an analytical solution of $u_1 = u_2 = \sin(\pi x)\sin(\pi y)$, and the right hand side is computed according to (1)-(2). The initial guess for the Newton iterations is a uniform displacement of 0, and we then solve (3) with the given right hand side and boundary conditions. To test for scalability, the domain is partitioned into $2 \times 2, 4 \times 4$ and 8×8 subdomains, respectively.

The computational results are reported in table 2. First consider the case of a splitting into 2×2 subdomains. As expected, non-linear preconditioning reduces the number of global Newton iterations, and the non-linear coarse component further improves the results. In terms of global linear iterations, linear DD and the 2-level ASPIN method have about the same number of iterations per Newton iteration, while the 1-level ASPIN method suffers from the lack of a coarse component in (8). The total number of local solves show that non-linear preconditioning clearly is worthwhile in this case.

Next, consider the performance as the number of subdomains is increased. Both the non-linear and the linear 2-level methods scale well: The number of global Newton iterations is more or less constant, and the number of linear iterations decreases somewhat. The 1-level

Table 2: The number of global Newton (GN), local Newton (LN), global GMRES (GL) and local linear solves (LL) for the homogeneous test problem. The local linear solves are composed by the local Newton iterations and the preconditioning of the global linear iterations

Sub-domains	2x2				4x4				8x8			
	GN	LN	GL	LL	GN	LN	GL	LL	GN	LN	GL	LL
ASPIN 2	4	48	81	372	4	135	69	1239	4	468	54	3942
ASPIN 1	6	56	181	780	7	224	294	4928	9	1089	516	34113
LIN-DD	8	0	173	692	7	0	122	1952	7	0	86	5504

ASPIN method sees a significant increase in the number of global Newton iterations, and for the finest partitioning it is outperformed by the standard Newton method. The non-local character of the non-linearities means that, unless combined with a coarse solver, attempts at localizing the computations will suffer from inaccurate boundary conditions, and the approximation deteriorates as the local boundary conditions see less of the global state. This trend continuous for smaller subdomains finer than those reported here.

4.2 Material heterogeneities

The second test case resembles a composite bar where circular inclusions are surrounded by a frame, see figure 1. The domain is $\Omega = [0,20] \times [0,1]$, and it is split into 20 subdomains of unit size. There is one circular inclusion in each subdomain, again see figure 1. The stress tensor is again given by (2). The inclusions has a coefficient $\alpha_i = 10^9$, while for the frame, we vary α_f between 10^4 and 10^9. Define $\gamma = \log_{10}(\alpha_i/\alpha_f)$. The remaining parameters in (2) are homogenous, and set to $\mu = 10^3, \lambda = 10^2$ and $\beta = 3$. The displacement is fixed to 0 at $x = 0$, and zero Neumann conditions are assigned to the rest of the boundary. The bar then bends under its own weight, rendering a non-linear elasticity problem. As initial guess for the iterations, we solve the linearized version of (1)-(2), that is, set $\alpha = 0$.

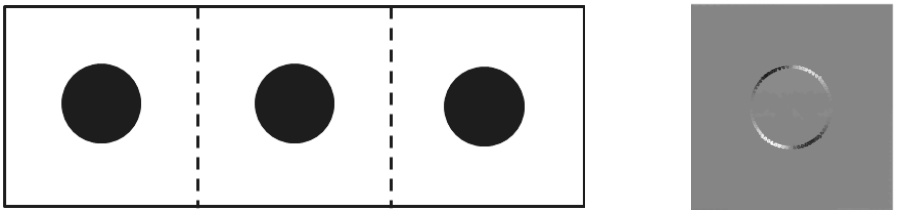

Figure 1: Left: Conceptual drawing of the computational domain used in section 4.2. The parameter $\boldsymbol{\alpha}$ controlling non-linearity is fixed in the black circles, and reduced in the surrounding frame. Stippled lines indicate subdomain boundaries. Right: Localization of the residual after the coarse-scale solve for a high contrast $\boldsymbol{\gamma}$; away from the material discontinuity the residual is essentially zero.

Table 3: The number of global Newton (GN), local Newton (LN), global GMRES (GL) and local linear solves (LL) for the heterogeneous test problem. The local linear solves are composed by the local Newton iterations and the preconditioning of the global linear iterations

γ	0				1				2			
	GN	LN	GL	LL	GN	LN	GL	LL	GN	LN	GL	LL
ASPIN 2	5	292	60	1492	4	182	32	822	4	154	27	694
ASPIN 1	8	354	1082	21994	6	226	853	17286	4	159	589	11839
LIN-DD	8	0	73	1492	7	0	56	1120	6	0	40	800
γ	3				4				3			
	GN	LN	GL	LL	GN	LN	GL	LL	GN	LN	GL	LL
ASPIN 2	4	158	32	798	4	160	28	720	4	160	28	720
ASPIN 1	3	130	453	9190	3	126	457	9266	3	126	457	9266
LIN-DD	6	0	40	800	6	0	41	820	6	0	41	820

The performance of the three non-linear solvers is listed in table 3. When the material properties are homogeneous ($\gamma = 0$), the 2-level method has the lowest number of global Newton iterations, whereas the 1-level ASPIN and standard Newton has comparable performance in this case. As γ increases, both the 2-level ASPIN method and standard Newton method behave somewhat better, with a relative difference in performance that is quite stable. However, the 1-level method improves significantly with increasing γ, and for $\gamma \geq 3$ it outperforms the 2-level ASPIN method both in terms of global Newton iterations, and in local non-linear iterations. As γ increases, the frame in figure 1 becomes more linear, thus the enhanced performance of the 1-level method is caused by the localization of non-linearities within each subdomain. In terms of global linear iterations, the 2-level ASPIN method outperforms the standard Newton method, whereas the 1-level method naturally suffers from the lack of a coarse preconditioner, independent of the contrast ratio γ.

To understand the comparatively weak results of the 2-level method, figure 1 also shows the residual in the subdomain with the highest non-linearity (the leftmost) after the coarse non-linear solve in the first global Newton iteration when $\gamma = 5$. The figure shows that the residual is confined to the boundary of the highly non-linear inclusion. Since the coarse-scale operators R_0 and P_0 are not formulated to adapt to the fine-scale state, they cannot accurately project the coarse-scale update to the fine-scale. As a result, the initial guess for the local fine-scale solves in the 2-level method are unfavorable compared to the 1-level method, leaving more of the burden to the global Newton iterations.

5 CONCLUDING REMARKS

In this paper we have discussed a two-level formulation of a non-linear preconditioner, formulated in the ASPIN framework. A non-linear coarse-scale problem was derived from the fine-scale state, and the coarse-scale update serves as initial guess for local fine-scale solves. Numerical examples from non-linear elasticity benchmarked the 2-level ASPIN method with a 1-level method, and a standard Newton solver combined with an inner GMRES solver that was preconditioned with a linear 2-level DD solver.

Numerical experiments show that for problems where the non-linearities do not localize, the non-linear coarse solver reduced the number of global Newton iterations compared to both 1-level ASPIN and standard Newton. Moreover, the coarse component renders scalability for the GMRES solver for an increasing number of subdomains. For problems with heterogeneous material fields, where the non-linearities do linearize, the coarse solver still give scalability for GMRES. However, the coarse restriction and prolongation operators failed to respect the fine-scale state, and the 1-level ASPIN method outperformed the 2-level method in terms of the number of non-linear solves in this case. Thus for the 2-level method to be competitive for general problems, more sophisticated restriction and prolongations are needed.

It should be mentioned that for the somewhat small problems considered herein, solving coarse non-linear problems are relatively cheap, and the cost has therefore been neglected. However, for large-scale and massively parallel computations with many sub-domains, solving the coarse problem becomes a bottleneck. The two natural solutions are either to only employ a linear coarse model [10], and thus keep the first benefit but loose the second, or to introduce a further coarsening step to arrive at a 3-level method. The simulations presented herein clearly show the merit of resolving non-linearities on the coarse scale in some cases, but more work is needed to identify efficient and robust solvers for large-scale non-linear problems.

REFERENCES

[1] M. Dryja and W. Hackbusch, On the nonlinear domain decomposition method, *BIT* (1997), **37:** 296–311.

[2] X.-C. Cai and D. E. Keyes, Nonlinearly Preconditioned Inexact Newton Algorithms, *SIAM J. Sci. Comput.*, (2002) **24**:183–200.

[3] N. Kikuchi and J. T. Oden, *Contact Problems in Elasticity: A Study of Variational Inequalities and Finite Element Methods*. SIAM, 1988.

[4] X.-C. Cai, M. Dryja, and M. Sarkis, Restricted additive Schwarz preconditioners with harmonic overlap for symmetric positive definite linear systems, *SIAM J. Numer. Anal.*(2003), **41**: 1209–1231.

[5] H.-B. An, On Convergence of the Additive Schwarz Preconditioned Inexact Newton Method, *SIAM J. Numer. Anal.* (2005), **43**: 1850–1871.

[6] J. E. Dennis and R. B. Schnabel, *Numerical Methods for Unconstrained Optimization and Nonlinear Equations*. SIAM, 1983.

[7] L. Marcinkowski and X. Cai, Parallel performance of some two-level ASPIN algorithms, in *Domain Decomposition Methods in Science and Engineering*, 2005, pp. 639–646.

[8] X.-C. Cai and M. Sarkis, A restricted additive Schwarz method for general sparse linear systems, *SIAM J. Sci. Comput.* (1999), **21**: 792–797.

[9] Sintef, Matlab Reservoir Simulation Toolbox, *http://www.sintef.no/Projectweb/MRST/*, 2013.

[10] F.-N. Hwang and X.-C. Cai, A class of parallel two-level nonlinear Schwarz preconditioned inexact Newton algorithms, *Comput. Methods Appl. Mech. Eng.* (2007) 196: 1603–1611.