

Interactive comment on “Pragmatic Solvers for 3D Stokes and Elasticity Problems with Heterogeneous Coefficients: Evaluating Modern Incomplete LDL^T Preconditioners” by Patrick Sanan et al.

Marcin Dabrowski (Referee)

mdabr@pgi.gov.pl

Received and published: 25 June 2020

General comments

The paper addresses problems relevant to the scope of SE and includes some interesting novel concepts regarding the numerical solution of 3D mechanical problems. The scientific methods and assumptions are sound and the presented conclusions are justified. The authors give credit to previous related work and they clearly delineated their contribution. The paper is well written and properly structured, the title is informative

Printer-friendly version

Discussion paper



and clear, and the abstract provides a good summary of the work.

Below I present my specific comments and technical corrections. I would encourage the author to include a more detailed presentation of the studied numerical setups within the main body of the manuscript, according to my detailed suggestions below.

Specific comments

For the Taylor-Hood element, the static elasticity in the mixed finite element formulation produces a symmetric indefinite system. It is maybe worth noting that for FEM discretization with piecewise discontinuous pressure field such as in the case of the Crouzeix-Raviart element family, the pressure mass matrix can be easily inverted on the element level and by performing block Gaussian elimination a positive definite system can be obtained that allows for using the highly robust sparse Cholesky factorization.

The author claimed that the previous incarnations of ILDL' were not necessarily robust (I.93-94). Could the authors just briefly mentioned the major improvements within the recent ILDL' implementation? What improvements exactly have made them robust in the recent years?

What is exactly meant by "coefficient structure" (for example I. 115)? I guess that this is not just the sparsity pattern.

1x1 and 2x2 blocks are mentioned in the context of pivoting for both LDL' and ILDL'. I am actually wondering whether the natural blocking inherent to the problem due to its dimensionality is retained during this operation? It is stated that fill-reducing reordering is performed block-wise. Which blocks are exactly meant here? I would guess that the ones related to the problem dimensional (say 3x3 blocks in the case of 3D problems) How is it ensure that the blocking due to the symmetric maximum weighted matching preprocessing is retained during the subsequent fill-in reducing reordering? I would suggest that this issue could be clarified in the manuscript.

[Printer-friendly version](#)[Discussion paper](#)

So what is exactly used as the Schur complement preconditioner for large coefficient jumps? The authors mention “a scaled pressure mass matrix” in this context. What (viscosity) scaling is exactly used? If there is not enough space for explaining it, maybe the authors could refer to some other work here.

The authors claim that sparse direct solution methods for indefinite systems using LDL' are expected to be highly competitive for 2D cases (l. 263). Is there any recent study showing their real performance (not just the theoretical scaling) that could be referred here?

The authors make statements that variable coefficients on the level of individual elements (non-grid aligned coefficient jumps) are, loosely speaking, harder to solve. In what sense? Solution accuracy or solution time, or maybe both?

I would guess that referring to incomplete factorization preconditioners in l. 283, the authors specifically mean ILDL' rather than ILU or ICHOL, and that they perhaps make this statement in the context of geodynamics or, in general, geosciences.

The comparative study of Gould (2007) is mentioned in l. 292 to justify the choice of PARDISO. I am wondering where there could be any more up-to-date performance studies for sparse direct solvers of symmetric indefinite systems.

I totally agree with the authors that the choice of norm is important for matrices characterized by large condition numbers such as in the case of the studied Stokes problem with strongly variable viscosities. In this respect, the authors choose to use the true residual 2-norm rather than the norm induced by the preconditioning. It is perhaps outside the scope of this study, but, in my view, given that the authors have access to highly accurate solutions obtained using the direct solver approach, it would be quite interesting to check and compare the solution error between GMRES(60)/ILDL & FGMRES(30)/ABF, say in the energy norm.

Do the solve times reported in the tables for the ILDL' preconditioning include the time

[Printer-friendly version](#)[Discussion paper](#)

spent on computing the ILDL' preconditioner? Actually, it would interesting to see how this time compares to the time spent on iterations.

Regarding the numerical setup, I would claim that what really matters is the fraction of the inclusion. With increasing inclusion fraction, as in the case of the setup studied in Fig. 2, a natural transition towards porous media like systems occurs (technically speaking, I am wondering how well 100 inclusions can be resolved using a 32^3 computational mesh). Such physical systems are characterized by strongly localized flows, which might be harder to solve compared to the suspension type of flow typically obtained for low concentration. It would be actually interesting to see how well the presented methods work when the gravity load is replaced by an ambient pressure gradient prescribed through boundary traction.

Technical corrections

I. 11-14 This sentence seems a bit convoluted. I would actually guess that something might be missing here.

I. 22-23 ... the coefficient structure is made increasingly challenging – I would suggest formulating it more precisely; What “complex topologies: have been addressed in this study?

I. 33 This is maybe not so critical, but compressible quasi-static linear elasticity is not exactly an example of a problem with a divergence free displacement field. In addition, it may indeed represent a saddle point problem, but in some numerical formulations it may be straightforwardly cast as a positive-definite problem.

I.71 ... the nonzero entries of the factors are restricted to those for A^k – This could be stated a bit more precisely.

I.72-73 I find the end part of this sentence unclear.

I. 109 In contrast to the previous ILDL studies previous mentioned above ... - please remove the second instance of “previous”

Printer-friendly version

Discussion paper



I. 134 One could consider using the transpose for one of the vectors in $n^* \sigma^* n$, etc in eq. 3 & 4.

I. 179 Is it really necessary to replace τ with $\text{dev}(\sigma)$ in eq. 3? Given that the t and n vectors are perpendicular ($\langle t, n \rangle = 0$), $t^* \sigma^* n = t^* (\tau - p \cdot I) \cdot n = t^* \tau^* n - p \cdot t^* n = t^* \tau^* n$

I. 210 permutation (a map from rows to columns) – I would think that the permutation operates within the rows and within the columns, and not from rows to column.

I. 215 If one wishes to find a symmetric permutation, one can only change the order of the diagonal entries. – If I am getting it right, a symmetric permutation preserves the symmetry of the matrix. I guess that with changing the order of the diagonal entries, the order of the entire rows and columns is also changes (not just the order of the diagonal entries). Anyway, could “non-symmetric” permutations be considered in the current context?

I. 293-4 Through a custom interface we use PARDISO (Kuzmin et al., 2013) – This looks a bit repetitive with respect to the previous sentence.

I. 298 The choice or norm allows is . . . - Please fix.

Table 1 - I would suggest that the volume fraction of the inclusions could be given. The viscosity is shown without the unit, and this problem could be easily solved by showing the viscosity ratio. Is the relative density dimensionless? Is it defined as $(\rho_{\text{incl}} - \rho_{\text{host}}) / \rho_{\text{host}}$? Is it actually relevant given that the model is linear? I would suspect that changing the relative density should only result in a rescaled velocity, and it should, hopefully, produce no appreciable changes to the course of numerical iterations. Is “fill” defined as the ratio between the non-zero entries in the LDL’ factor with respect to the non-zero entries of the original matrix (the triangular part of it, including the diagonal)? Is it necessary to use the scientific notation when time is reported? Maybe giving the total dof count could be useful.

Printer-friendly version

Discussion paper



Figure 1 - I would suggest a more detailed description of the numerical setups, both in the caption and in the main body of the manuscript. What is the volume fraction of the inclusions? What is meant by (Vel. scaled $1/3x$)? Isn't it that the scaling of the quiver lengths is in no obvious way absolute? I think that it would be useful to show grid lines in the plots. I guess that the dashed line in the Peak Memory Footprint shows the maximum available RAM during the numerical tests, but it would be useful to explain it in the caption. The curve styles are not well visible in the legend. It could also be explicitly explained in the caption that ABF(a), ABF(b), ... refer to setups a, b, c... (at a first glance it may look as if it were some variants of the solvers).

I. 317-8 ... the ABF solver fails to converge. – It is not clear to me where this can be seen in Fig. 1 (I can't really see any missing data for ABF)

Figure 2 - What is the volume fraction of the inclusions as their number is increase? Given that the numerical resolution is kept constant (32^3) I would guess that it is increased. In my opinion, this should be explicitly stated in the caption and also in the main body of the manuscript. In fig. 1 for 32^3 the overall solver performance in terms of dof/s fell in to the range between $5 \cdot 10^3$ and 10^4 , which is consistent with the time reported in table 1. However, in fig. 2, even in the previously studied case of the viscosity ratio of 10^4 , the performance is between 10^{-2} and 10^{-1} . I would guess that this could be some technical mistake. In my opinion, it would be useful to show gridlines and maybe use a slightly large font for the legend entries.

I.324 "... varying to drop tolerance" – Please fix.

I. 326 System scaling is mentioned in the footnote. Please explain what system scaling (physical, algebraic, ..) is exactly meant here.

I.339.. and C is the term (depending on λ as in Eq.(9)). – I would guess that the outer brackets are not necessary here. Could the author hint what they actually use for the C term?

[Printer-friendly version](#)[Discussion paper](#)

I. 340 Figure 4 shows a similar experiment using a scenario which is perhaps more typical in applications. – Please explain the boundary conditions used in this setup in the main body of the manuscript.

Figure 3 – Maybe the Lamé parameters μ and λ could be scaled by $\rho \cdot g \cdot L$. A colorbar for the color-coded pressure and gridlines would be a nice addition to this figure.

Figure 4 – It is of small relevance to the studied topic, but the deformed wire mesh implies a substantial deformation that could hardly be accommodated elastically by any geomaterial. But maybe this could be treated as an exaggerated mesh deformation. The elastic moduli are given with no units.

Marcin Dabrowski

Interactive comment on Solid Earth Discuss., <https://doi.org/10.5194/se-2020-79>, 2020.

Printer-friendly version

Discussion paper

