

Advantages of static condensation in implicit compressible Navier-Stokes DGSEM solvers

Wojciech Laskowski^a, Andrés M. Rueda-Ramírez^{ac}, Gonzalo Rubio^{ab},
Eusebio Valero^{ab}, Esteban Ferrer^{ab}

^a*ETSIAE-UPM (School of Aeronautics - Universidad Politécnica de Madrid) - Plaza de Cardenal Cisneros 3, 28040 Madrid, Spain*

^b*Center for Computational Simulation - Universidad Politécnica de Madrid, Campus de Montegancedo, Boadilla del Monte, 28660 Madrid, Spain*

^c*Department for Mathematics and Computer Science, University of Cologne, Weyertal 86-90, 50931, Cologne, Germany*

Abstract

We consider implicit time-marching schemes for the compressible Navier-Stokes equations, discretised using the Discontinuous Galerkin Spectral Element Method with Gauss-Lobatto nodal points (GL-DGSEM). We compare classic implicit strategies for the full Jacobian system to our recently developed static condensation technique for GL-DGSEM *Rueda-Ramírez et al. (2019), A Statically Condensed Discontinuous Galerkin Spectral Element Method on Gauss-Lobatto Nodes for the Compressible Navier-Stokes Equations* [1]. The Navier-Stokes system is linearised using a Newton-Raphson method and solved using an iterative preconditioned-GMRES solver. Both the full and statically condensed systems benefit from a Block-Jacobi preconditioner.

We include theoretical estimates for the various costs involved (i.e. calculation of full and condensed Jacobians, factorising and inverting the preconditioners, GMRES steps and overall costs) to clarify the advantages of using

static condensation in GL-DGSEM, for varying polynomial orders. These estimates are then examined for a steady three-dimensional manufactured solution problem and for an two-dimensional unsteady laminar flow over a NACA0012 airfoil. In all cases, we test the schemes for high polynomial orders, which range from 2 to 8 for a manufactured solution case and from 2 to 5 for the NACA0012 airfoil. The statically condensed system shows computational savings, which relate to the smaller system size and cheaper Block-Jacobi preconditioner with smaller blocks and better polynomial scaling, when compared to the preconditioned full Jacobian system (not condensed). The advantage of using static condensation is more noticeable for higher polynomial orders.

Keywords:

High-order discontinuous Galerkin, DGSEM, Gauss-Lobatto, Implicit time-marching, preconditioned-GMRES, Compressible Navier-Stokes, Static condensation, NACA0012 airfoil

Contents

1	Introduction	3
2	Methodology	8
2.1	Time-implicit discretisation and Jacobian computation	9
2.2	Static condensation	11
2.3	Size of the full and the condensed Jacobians	12
2.4	Preconditioned-GMRES solver	16
2.5	Further implementation details	17

3	Theoretical costs of full and statically condensed systems	18
3.1	Cost of static condensation	20
3.2	Cost of factorising the preconditioner	22
3.3	Cost of the preconditioned-GMRES solver	23
3.4	Summary of computational costs	25
4	Numerical results	27
4.1	Steady simulation: Manufactured Solution	28
4.2	Unsteady simulation: NACA0012 at AOA = 20°	33
5	Conclusion	39
Appendix A	Preliminary assessment of preconditioners	42
Appendix B	Influence of Mach and Reynolds	45
Appendix C	Estimation of non-zero entries in the Jacobian Matrix	46
Appendix C.1	Advection terms	47
Appendix C.2	Diffusion terms	50
Appendix C.3	Total number of non-zero entries	55

1 **1. Introduction**

2 The accurate simulation of aerodynamic characteristics over lifting sur-
3 faces (airfoils and wings) is of major importance to the aeronautical industry
4 and can potentially reduce fuel consumption by allowing lighter aircraft de-
5 signs. High order methods, and particularly discontinuous Galerkin (DG)

6 schemes, are well equipped to provide high accuracy on coarse meshes due
7 to their spectral convergence property (i.e. exponential decay of the error).
8 In the last decade, these methods have gained popularity for solving fluid
9 flows governed by the incompressible, e.g. [2, 3, 4, 5, 6] and compressible
10 Navier-Stokes equations, e.g. [7, 8, 9, 10]. DG solutions show improved ac-
11 curacy over low order methods, but are often expensive to compute [11]. In
12 recent years, acceleration techniques for DG schemes have focused on local
13 p -adaption, see e.g. [7, 10] and on improved time-marching techniques, e.g.
14 FAS p -multigrid [10], that allow for faster convergence and large time-steps,
15 with important savings in computational cost.

16 The Discontinuous Galerkin Spectral Element Method (DGSEM) [12],
17 is a particular nodal version of DG, which has proved to be very efficient
18 on hexahedral elements (e.g. diagonal mass matrices). Additionally, the
19 variant of the DGSEM where Gauss-Lobatto nodal points are selected, i.e.
20 GL-DGSEM, is well suited for the development of provably stable schemes
21 [13], fulfilling the summation-by-parts property [14]. These schemes have
22 enhanced stability and are convenient for under-resolved simulations, if split-
23 forms of the governing equations are discretised. Examples of provably stable
24 formulations can be found for the Euler [14], the Magneto-Hydrodynamics
25 [15], multiphase flows [16, 17] and the Navier-Stokes equations [18, 9, 19].

26 We have recently shown an additional advantage of GL-DGSEM [1]: it
27 is well suited for the static condensation approach, whilst the classic Gauss
28 point version is not. In this work we exploit the statically condensed system,

29 to accelerate implicit time advancement with an iterative GMRES solver,
30 and compare the accelerations to the traditional full Jacobian system. Note
31 that both approaches rely on Newton-Raphson linearisation to obtain the
32 full and condensed systems. In this work, we do not include split-forms
33 but propose a static condensation technique, which is perfectly applicable
34 to formulations including stabilising split-forms (e.g. two point fluxes), and
35 may be combined with the static condensation, in future work.

36 Static condensation has been widely applied in the context of high order
37 methods, and is a popular strategy in the continuous Galerkin community,
38 e.g. [20, 21], where it has proved to be an efficient strategy to solve large
39 systems in both structural and fluid mechanics, e.g. [20, 22]. Static conden-
40 sation can be combined with modern iterative techniques such as p -multigrid
41 with domain decomposition smoothers tailored for condensed systems [23].
42 Recently, Pardo et al. [24] showed that static condensation proves beneficial
43 when combined with iterative solvers, if the number of iterations is suffi-
44 ciently large, to compensate for the additional cost associated of computing
45 the system's Schur complement. Similar findings are included in this work
46 for DGSEM.

47 Static condensation has been applied to discontinuous Galerkin discreti-
48 sations by Sherwin et al. [25] and Hybridized Discontinuous Galerkin (HDG),
49 e.g. [26, 27, 28]. In the first work, Sherwin et al. reported advantages of stat-
50 ically condensed systems when using tailored non-orthogonal basis functions
51 (i.e. non-diagonal mass matrices). The remaining references were developed

52 for HDG formulations, where the method decouples the degrees of freedom
53 belonging to the mesh elements from the mesh skeleton, enabling static con-
54 densation. However, HDG requires specific numerical fluxes [1, 26, 29], re-
55 stricting the use of well known Riemann approximations, such as Roe’s. Our
56 static condensation for GL-DGSEM allows any flux.

57 In our previous work [1], we showed the detailed implementation of the
58 static condensation approach in GL-DGSEM, and applied the method to
59 solve steady cases using direct solvers and an implicit GMRES with a point-
60 Jacobi preconditioner. In this work, we extend that analysis further by
61 comparing the performance of statically condensed and full Jacobian (non-
62 condensed) systems for Block-Jacobi preconditioner in steady and unsteady
63 problems, and show that the statically condensed system can lead to faster
64 iterative GMRES solves. We include theoretical estimates to analyse and
65 extrapolate the costs involved with respect to the polynomial order. These
66 include the calculation of full and condensed Jacobians, the factorisation and
67 inversion of the preconditioner and the preconditioned-GMRES steps. Ad-
68 ditionally, we briefly asses the use of ILU(k) preconditioners and include a
69 section to verify that the advantages of the statically condensed GL-DGSEM
70 are essentially independent of the Mach and Reynolds numbers.

71 Both full and condensed systems can benefit from preconditioners to
72 accelerate convergence. Efficient preconditioners should be cheap to con-
73 struct and to parallelise, whilst enhancing the convergence of the system,
74 e.g. reducing the number of iterations to reach convergence. Iterative strate-

gies (including preconditioners) for DG discretisations of both compressible and incompressible flows have been widely explored in recent years [30, 31, 32, 33, 34, 35, 36, 37, 28, 38, 39]. Most authors employ block structured preconditioners/ p -multigrid smoothers, such as Block-Jacobi, Line-Jacobi, additive-Schwarz or Block-ILU. Among these, [37, 28, 39] focused on coarse grid accelerations and efficient implementation of the *state-of-the-art* solvers for turbulent problems, which is out of the scope of this work. Point ILU has also been successfully used for aerodynamic applications in [40, 41]. Persson and Peraire [32] or Gopalakrishnan and Kanschat [42] showed that element-block based preconditioners are essential to eliminate high p dependent errors. It is also very natural to exploit the element-block structure of the Jacobian (specially in the parallel computations due to the block locality that enables to perform block inversions locally), as most of these methods require the direct factorisation of block matrices. Note that this can become troublesome for high polynomial orders, especially in three-dimensional flows. In this work, we select Block-Jacobi preconditioner and show that when condensing the system, the preconditioner scales more gently for high polynomials, than the preconditioner for the full system. This translates into lower costs for all the steps where the preconditioner is required (i.e. factorisation of the blocks and GMRES step involving the preconditioner), and paves the way to using high polynomial orders efficiently.

Our comparisons are novel in that the static condensation technique, recently developed for GL-DGSEM by the authors, is directly challenged to

98 the state of the art implicit preconditioned-GMRES solvers to show com-
99 putational savings for steady and unsteady flows and a range of polynomial
100 order ranging from 2 to 8. The results are backed-up by the theoretical es-
101 timates for the various costs. The beneficial effect of statically condense the
102 system is observed for various Mach and Reynolds numbers, suggesting that
103 this technique can be exploited for a wide range of flow regimes in steady
104 and unsteady flows.

105 In what follows, we describe the methodology with emphasis on the time
106 marching scheme and implementation details. We continue with the theo-
107 retical estimates and the simulations, where we compare the full Jacobian
108 and the static condensation for a 3D Manufactured Solution problem and
109 the unsteady flow over a 2D NACA0012 airfoil. We finalise with conclusions
110 and outlooks.

111 **2. Methodology**

112 We use the nodal Discontinuous Galerkin Spectral Element Method (DGSEM)
113 introduced by Black [43], where the computational domain is tessellated into
114 non-overlapping hexahedral elements. In the DGSEM, numerical fluxes are
115 necessary to transfer information between discontinuous element solutions.
116 Here, we retain Lax-Friedrichs fluxes for the convective fluxes and the Inte-
117 rior Penalty method for viscous fluxes, but other fluxes with compact support
118 could also be used (e.g. Roe for convection or BR2 for diffusion). The se-
119 lected fluxes yield a compact mesh stencil and are differentiated to obtain

120 an analytical Jacobian. Further details on how the Jacobian can be ob-
 121 tained along with the peculiarities and sparsity patterns resulting from using
 122 Gauss-Lobatto nodal points, can be found in our previous works [1, 44].

123 *2.1. Time-implicit discretisation and Jacobian computation*

124 Let us briefly describe the implicit methods retained in this work. After
 125 discretising the compressible Navier-Stokes equations, we obtain the follow-
 126 ing system of equations

$$\underline{\mathbf{M}} \frac{\partial \mathbf{Q}}{\partial t} + \mathbf{F}(\mathbf{Q}) = \underline{\mathbf{M}} \mathbf{S}, \quad (1)$$

127 where \mathbf{Q} is a vector that stores the conservative variables in all degrees
 128 of freedom of the domain, $\mathbf{F}(\mathbf{Q})$ encompasses both discrete convective and
 129 diffusive fluxes, $\underline{\mathbf{M}}$ is the mass matrix, which is diagonal in the nodal DGSEM
 130 approach, and \mathbf{S} is a source term.

We replace the continuous in time derivative in (1) by a discrete implicit
 time integration scheme using Backward Differentiation Formulas (BDF) of
 order 1 and 2 (BDF1 or BDF2),

$$\frac{\partial \mathbf{Q}}{\partial t} \leftarrow \frac{\delta \mathbf{Q}}{\delta t}(\mathbf{Q}_{s+1}, \mathbf{Q}_s, \dots), \quad (2)$$

where the operator $\delta \mathbf{Q} / \delta t$ is a function of the solution on the next time step,
 \mathbf{Q}_{s+1} (the unknown), the current time step, \mathbf{Q}_s , and possibly previous time
 steps. When treated implicitly, the nonlinear operator \mathbf{F} , in equation (1) is

evaluated for the unknown solutions, \mathbf{Q}_{s+1} . Considering this, equation (1) can then be rewritten as

$$\mathbf{R}(\mathbf{Q}_{s+1}) = \frac{\delta \mathbf{Q}}{\delta t}(\mathbf{Q}_{s+1}, \mathbf{Q}_s, \dots) + \underline{\mathbf{M}}^{-1} \mathbf{F}(\mathbf{Q}_{s+1}) - \mathbf{S} = \mathbf{0}. \quad (3)$$

131 Note that in the DGSEM approach the mass matrix $\underline{\mathbf{M}}$ is diagonal and can
 132 be trivially inverted, leading to an efficient discontinuous Galerkin method.
 133 When computing steady flows, we are not interested in producing an accurate
 134 solution in time, and therefore we use an implicit BDF of order 1 to advance
 135 until steady state. However, for unsteady cases we will use an implicit BDF
 136 of order 2 and shorter time steps to obtain accurate solutions in time.

137 The nonlinear system of equations, (3), can be solved using Newton-
 138 Raphson iterations to obtain the linear system:

$$\underline{\mathbf{A}} \Delta \mathbf{Q} = \mathbf{B}, \quad (4)$$

139 where $\underline{\mathbf{A}} = \frac{\partial \mathbf{R}}{\partial \mathbf{Q}}(\tilde{\mathbf{Q}}_{s+1})$ is the Jacobian matrix evaluated at $\tilde{\mathbf{Q}}_{s+1}$, which is
 140 an approximation to the unknown solution \mathbf{Q}_{s+1} . The right-hand-side is
 141 $\mathbf{B} = -\mathbf{R}(\tilde{\mathbf{Q}}_{s+1})$. Equation (4) is a linear system that must be solved
 142 iteratively to approach $\mathbf{Q}_{s+1} \leftarrow \tilde{\mathbf{Q}}_{s+1} + \Delta \mathbf{Q}$. The Jacobian matrix $\underline{\mathbf{A}}$ may
 143 be computed analytically or numerically, and here we retain the analytical
 144 approach, for its efficiency. Equation (4) is what we refer as *full system* with
 145 $\underline{\mathbf{A}}$ the *full Jacobian*.

146 *2.2. Static condensation*

147 In the GL-DGSEM framework, we can statically condense system (4) to
 148 obtain the following form

$$\begin{bmatrix} \underline{\mathbf{A}}_{bb} - \underline{\mathbf{A}}_{ib}\underline{\mathbf{A}}_{ii}^{-1}\underline{\mathbf{A}}_{bi} & \mathbf{0} \\ \underline{\mathbf{A}}_{bi} & \underline{\mathbf{A}}_{ii} \end{bmatrix} \begin{bmatrix} \Delta\mathbf{Q}_b \\ \Delta\mathbf{Q}_i \end{bmatrix} = \begin{bmatrix} \mathbf{B}_b - \underline{\mathbf{A}}_{ib}\underline{\mathbf{A}}_{ii}^{-1}\mathbf{B}_i \\ \mathbf{B}_i \end{bmatrix}, \quad (5)$$

149 where subindex b and i denote boundary and interior nodes, respectively.
 150 The main interest of the method is to obtain a block diagonal matrix $\underline{\mathbf{A}}_{ii}$,
 151 that can be inverted cheaply and locally (element by element). Additionally,
 152 the boundary matrix including the degrees of freedom linking boundaries
 153 between elements, is greatly reduced by the use of Gauss-Lobatto points in
 154 DGSEM [1]. The resulting system is equivalent to the full system, but can
 155 be decoupled in two subsystems. The first one for the skeleton of the mesh,
 156 our *condensed system of equations* is

$$\underline{\mathbf{A}}_{cond}\Delta\mathbf{Q}_b = \mathbf{B}_{cond}, \quad (6)$$

157 where $\underline{\mathbf{A}}_{cond} = \underline{\mathbf{A}}_{bb} - \underline{\mathbf{A}}_{ib}\underline{\mathbf{A}}_{ii}^{-1}\underline{\mathbf{A}}_{bi}$ and $\mathbf{B}_{cond} = \mathbf{B}_b - \underline{\mathbf{A}}_{ib}\underline{\mathbf{A}}_{ii}^{-1}\mathbf{B}_i$. Once the
 158 condensed system (6), based on the Schur complement $\underline{\mathbf{A}}_{cond}$, is solved, then
 159 it is trivial to substitute and solve for the second system $\Delta\mathbf{Q}_i = \underline{\mathbf{A}}_{ii}^{-1}(\mathbf{B}_i -$
 160 $\underline{\mathbf{A}}_{bi}\Delta\mathbf{Q}_b)$, since $\underline{\mathbf{A}}_{ii}$ is block diagonal and has already being factorised to
 161 compute $\underline{\mathbf{A}}_{cond}$.

One of the main advantages of the static condensation is the reduced size of the matrix $\underline{\mathbf{A}}_{cond}$ (with only the mesh skeleton degrees of freedom) in comparison with the original Jacobian matrix $\underline{\mathbf{A}}$ (with all the degrees of freedom in the mesh). We can quantify the number of degrees of freedom for our GL-DGSEM discretisation. The Jacobian matrix $\underline{\mathbf{A}}$ has size

$$n = N_{el} \cdot nb. \quad (7)$$

where N_{el} is number of elements and nb is the size of each element-block. Then, assuming mesh elements with isotropic polynomial order P , we can describe the size of each block nb as a function of P , the dimension d (e.g. $d = 3$ for 3D meshes) and the number of conservative variables (or equations) in the computational domain for the Navier-Stokes equations N_{eq} (e.g. $N_{eq} = 5$ in 3D):

$$nb = N_{eq}(P + 1)^d. \quad (8)$$

Equation (7) can also be used to describe the size of the matrices, $\underline{\mathbf{A}}_{ii}$ and $\underline{\mathbf{A}}_{bb}$, involved in the Schur complement computation and included in the statically condensed system (5) with $n_{ii} = N_{el} \cdot nb_{ii}$ and $n_{bb} = N_{el} \cdot nb_{bb}$, with the only difference being the block sizes. Here, the block size of the element-skeleton matrix nb_{bb} directly corresponds to the size of the block of the final Schur complement $\underline{\mathbf{A}}_{cond}$. The blocks for the condensed matrix arise from having decoupled element interior i from the element boundary nodes

b , leaving fewer degrees of freedom per block. Thus, the size of the block of matrix $\underline{\mathbf{A}}_{ii}$, that corresponds to the interior of the elements is

$$nb_{ii} = N_{eq}(P - 1)^d. \quad (9)$$

Consequently, the size of the block of $\underline{\mathbf{A}}_{bb}$ and $\underline{\mathbf{A}}_{cond}$ can be defined as the difference between the size of the element-block and the interior element part

$$nb_{bb} = N_{eq} [(P + 1)^d - (P - 1)^d], \quad (10)$$

163 and with these blocks, the final size of the matrices could be easily computed
 164 from equation (7).

165 Additionally, it is possible to obtain estimates for the number of non-zero
 166 entries nnz in the full and condensed Jacobian. This is not a trivial task, and
 167 details are included in Appendix C. The final expressions are summarised in
 168 Table 1, for 3D and 2D.

Table 1: Explicit formulas for the leading terms of block sizes, estimation of number of non-zeros nnz per block, and matrix non-zero entries, for the full and condensed systems in 2D and 3D. All provided as functions of the number of elements N_{el} , polynomial order P and number of conservative variables in the 3D domain, i.e. $N_{eq} = 5$ for the compressible Navier-Stokes equations.

3D		
	Full system	Condensed system
Block size	$N_{eq}(P+1)^3$	$N_{eq}(6P^2+2)$
nnz per block	$3N_{eq}^2P(P+1)^4$	$N_{eq}^2(6P^2+2)^2$
nnz in matrix	$3N_{el}N_{eq}^2P(P+1)^4$	$25N_{el}N_{eq}^2(6P^2+2)^2$
2D		
	Full system	Condensed system
Block size	$N_{eq}(P+1)^2$	$N_{eq}4P$
nnz per block	$N_{eq}^2(P+1)^4$	$N_{eq}^216P^2$
nnz in matrix	$N_{el}N_{eq}^2(P+1)^4$	$13N_{el}N_{eq}^216P^2$

169 Let us remark that the expressions for the block sizes are exact. How-
170 ever, the expressions for the nnz per block are upper bounds derived in the
171 appendix. The entry corresponding to the nnz for the full system, only in-
172 cludes the diagonal blocks corresponding to the viscous terms, since these are
173 asymptotically dominant, as they scale $\mathcal{O}(P^5)$ (all other blocks have weaker
174 scaling, see appendix for details). The total number of non-zeros might be
175 obtained multiplying by the number of elements. Regarding the condensed
176 system, here the block stencil of this matrix is estimated to be 25 in 3D and
177 13 in 2D (neighbor to neighbor coupling), and therefore to obtain the total
178 number of nnz in the matrix, the nnz per block need to be multiplied by
179 the number of elements and by the constant (25 or 13) accounting for the
180 neighbour coupling.

181 Finally, the condensed system presents smaller and denser blocks and

182 the block stencil of the condensed system is wider than the one of the full
 183 system. As a result, the nnz of the condensed system is larger than the one
 184 of the full system. Regarding the total number of non-zero entries in the
 185 matrix, the scalings show that the full system will asymptotically contain
 186 more non-zero entries for large polynomial orders. However, due to the denser
 187 connectivity in the condensed system, the non-zero entries can be higher for
 188 low polynomial orders.

189 In the Continuous Galerkin formulation for simple diffusion or advection-
 190 diffusion problems [24, 45], the number of non-zero entries in the condensed
 191 matrix decreases with respect to nnz in the full system. However, in our
 192 case for the GL-DGSEM of the compressible Navier-Stokes equations, the
 193 number of non-zeros increases. Increased number of non-zeros for the con-
 194 densed system have been reported by Habchi [46], for an elasto-hydrodynamic
 195 lubrication problem. There, the authors considered several meshes for the
 196 same contact problem, from *extra coarse* to *extra fine*. The results show that
 197 nnz in the condensed systems is reduced for coarse meshes, whereas for the
 198 others $nnz_{cond} > nnz_{full}$.

199 Complementary illustrations of the static condensation sparsity patterns
 200 for the GL-DGSEM approach may be found in our previous work [1]. In this
 201 work, we concentrate on comparing the efficiency of solving the linear system
 202 of equations, i.e. solving full system (4) to solving the two subsystems for the
 203 condensed system (6) using iterative methods. To account for the iterative
 204 costs, we will use the matrix sizes and number of non-zeros, included in Table

205 1.

206 2.4. *Preconditioned-GMRES solver*

207 We use preconditioned-GMRES to solve both the full system (4), and the
208 statically condensed system (6). Previous works [33, 31, 32, 35, 36] have
209 shown that combining GMRES and block preconditioners is effective in solv-
210 ing Eq. (4) for DG discretisations of Euler, Navier-Stokes or RANS equations.
211 Here, we have considered several preconditioning strategies, namely element
212 Block-Jacobi and incomplete LU factorisation with different factorisation lev-
213 els, ILU(k). We conduct a preliminary evaluation of these preconditioners
214 for the full and condensed systems in Appendix A. For the manufactured
215 solution case (to be described later in detail), ILU(k) preconditioners per-
216 form better in terms of iteration count and overall cost, but show high cost
217 when computing the preconditioner. Block-Jacobi does not perform as well
218 as ILU(k) in terms of overall solver cost, but provides a lower factorisa-
219 tion cost (specially for the statically condensed system) and provides very
220 competitive average iteration count and average solver cost. Additionally, a
221 Block-Jacobi preconditioner is more suitable for parallel [28] and matrix-free
222 [47, 48] computations, since the blocks can be inverted locally whilst exploit-
223 ing the block-structure of the high order DGSEM discretisation, as well as
224 requiring less memory [33, 47]. For this reason, in the following sections, we
225 present all results with Block-Jacobi preconditioners for both the full (4) and
226 the statically condensed (6) systems.

227 The Block-Jacobi preconditioner ignores all the Jacobian off-diagonal
228 blocks and performs a local LU decomposition (factorisation step) in each
229 diagonal block. For the full system, these diagonal blocks include all the
230 element degrees of freedom for each element, whilst the size for the blocks is
231 reduced in the condensed system (only skeleton degrees of freedom): matrix
232 $\underline{\mathbf{A}}_{cond}$ in (6). These blocks are smaller as shown in Table 1 and therefore con-
233 structing the Block-Jacobi preconditioner for the condensed system is much
234 cheaper, than for the full system, and especially for high polynomial orders
235 (for a more detailed comparison of the factorisation costs, see next Section
236 3).

237 Finally, all the operations related to the preconditioned-GMRES solver
238 (computing preconditioner and performing GMRES iterations) are performed
239 using the well known open-source library PETSc [49, 50, 51]. The computa-
240 tion of the condensed system (6), however, is done with our in-house code.
241 Note that PETSc has been widely used in aeronautical publications, includ-
242 ing DGSEM flow simulations [33, 52, 53]. By selecting this well validated
243 implementation, we avoid in-house inefficiencies that could mask the out-
244 comes of our comparisons.

245 *2.5. Further implementation details*

246 In the result section, we also include explicit time-marching (ESRK3)
247 [54] simulations for reference, but comparisons of overall computing time
248 are not of interest in this work. Indeed, it is well known, that the explicit

249 time integrator is easy to parallelise with appropriate domain partitioning
 250 [55, 56] and could produce very efficient solutions when using large number
 251 of processors, whilst implicit schemes require a greater effort and increased
 252 memory requirements for matrix-based solvers [31, 11]. Alternative matrix-
 253 free approaches have been proposed, e.g. Pazner and Persson [48], but are
 254 not explored in this text. For the above mentioned reasons, all cases are run
 255 in serial such that all approaches are fairly compared without taking into
 256 account parallelisation strategies or communication efficiency.

257 **3. Theoretical costs of full and statically condensed systems**

258 In this section, a theoretical analysis of the main computational costs
 259 of the implicit time marching scheme are included. Algorithm 1 presents
 260 the essential steps of the time marching scheme to conduct the simulation
 261 until the finalisation criteria is met. We focus only in three main steps that
 262 constitute the majority of the computational costs, i.e.:

- 263 • Step 8: cost for obtaining the statically condensed system $\underline{\mathbf{A}}_{cond}$,
- 264 • Step 9: cost for factorising (constructing and inverting) the precon-
 265 ditioning matrix $\underline{\mathbf{P}}^{-1}$. In the context of this work, it is the cost of
 266 factorising the element-diagonal blocks of the Jacobian system $\underline{\mathbf{A}}$ or
 267 the condensed system $\underline{\mathbf{A}}_{cond}$, which are then inverted and stored in the
 268 preconditioning matrix $\underline{\mathbf{P}}^{-1}$.
- 269 • Step 13: cost for solving the linear system (4) for the full system or (6)

270 for the condensed system, using the preconditioned-GMRES solver at
 271 each time step and as long as $\|\Delta\mathbf{Q}\|_\infty < \text{TOL}_{Newton}$.

Algorithm 1 Time-marching scheme including Newton-Raphson linearisation

```

1:  $\mathbf{Q} \leftarrow \text{INITIALISE}()$ 
2: while Steady:  $\|\underline{\mathbf{M}}^{-1}\mathbf{F}(\mathbf{Q}) - \mathbf{S}\|_\infty < 10^{-8}$  or Unsteady:  $t < T_{end}$  do
3:    $t \leftarrow t + \Delta t$ 
4:   while  $\|\Delta\mathbf{Q}\|_\infty < \text{TOL}_{Newton}$  do
5:     if InaccurateJacobian then
6:        $\underline{\mathbf{A}} \leftarrow \text{COMPUTEFULLSYSTJACOBIAN}(\mathbf{Q}, \Delta t)$ 
7:       if CondensedSystem then
8:          $\underline{\mathbf{A}} \leftarrow \text{COMPUTECONDENSEDJACOBIAN}(\underline{\mathbf{A}})$ 
9:          $\underline{\mathbf{P}}^{-1} \leftarrow \text{FACTORISEPRECONDITIONER}(\underline{\mathbf{A}})$ 
10:       $\mathbf{B} \leftarrow \text{COMPUTEFULLSYSTRHS}(-\mathbf{R}(\mathbf{Q}))$ 
11:      if CondensedSystem then
12:         $\mathbf{B} \leftarrow \text{COMPUTECONDENSEDRHS}(\underline{\mathbf{A}}, \mathbf{B})$ 
13:       $\Delta\mathbf{Q} \leftarrow \text{GMRES-SOLVE}(\underline{\mathbf{A}}, \underline{\mathbf{P}}^{-1}, \mathbf{B})$ 
14:      if CondensedSystem then
15:         $\Delta\mathbf{Q} \leftarrow \text{COMPUTEINTERIORSOLUTION}(\underline{\mathbf{A}}, \Delta\mathbf{Q}, \mathbf{B})$ 
16:       $\mathbf{Q} = \mathbf{Q} + \Delta\mathbf{Q}$ 

```

272 Step 13 solves the linear system using preconditioned-GMRES (further
 273 discussed below) and one must account for its cost in every Newton iteration
 274 and for every time step. Steps 5 to 9 need to be computed when the Jaco-
 275 bian matrix $\underline{\mathbf{A}}(\mathbf{Q}, \Delta t)$ (or the condensed version), has significantly changed,
 276 which leads to a quasi-Newton method. Naturally, re-using the Jacobian
 277 matrix from the previous time steps may inhibit quadratic convergence of
 278 the Newton-Raphson method [57]. To ensure a sufficiently high convergence
 279 rate, we follow ideas from Zahr and Persson [58] and define a condition that

280 secures at least 1/4 of an order of magnitude decay per Newton iteration (see
 281 step 5 of Algorithm 1). Therefore, if the aforementioned condition is met, the
 282 Jacobian $\underline{\mathbf{A}}$ and preconditioner $\underline{\mathbf{P}}^{-1}$ are still useful and are not recomputed.
 283 In all the simulations, the Newton tolerance is set to $\text{TOL}_{\text{Newton}} = 10^{-5}$,
 284 which yields accurate results. Furthermore, as in Nastase and Mavriplis [59],
 285 the preconditioned-GMRES solver tolerance is set according to the maximum
 286 norm of the residual, e , such that $\text{TOL}_{\text{GMRES}} = e \cdot 0.7^i$, where i is the current
 287 Newton iteration.

288 Sections 3.1, 3.2 and 3.3 present the estimation of the computational costs
 289 related to the static condensation (Step 8), the preconditioner factorisation
 290 (Step 9) and the GMRES solver (Step 13). Subsequently, comparisons with
 291 the simulated costs are included in Section 4, and summarised in table 2.

292 3.1. Cost of static condensation

293 The necessary operations to obtain the condensed system (6) are detailed
 294 here:

- 295 • Factorisation and inverting the block diagonal matrix representing inner-
 296 element $\underline{\mathbf{A}}_{ii}^{-1}$,
- 297 • Computing $\underline{\mathbf{A}}_{ii}^{-1} \underline{\mathbf{A}}_{ib}$ and assembling the RHS of the equation (6),
- 298 • Computing the $\underline{\mathbf{A}}_{\text{cond}} = \underline{\mathbf{A}}_{bb} - \underline{\mathbf{A}}_{bi} \underline{\mathbf{A}}_{ii}^{-1} \underline{\mathbf{A}}_{ib}$, equation (6),
- 299 • Obtaining the solution for the interior nodes: $\Delta \mathbf{Q}_i = \underline{\mathbf{A}}_{ii}^{-1} (\mathbf{B}_i - \underline{\mathbf{A}}_{ib} \Delta \mathbf{Q}_b)$.

300 All of these operations are included in one unique cost, referred to as *con-*
 301 *densation cost*, in the following sections. These operations are performed in
 302 Step 8 in Algorithm 1. The only exceptions are obtaining the solution for the
 303 interior nodes, which is performed in step 15, and assembling the RHS of the
 304 equation (6), which is performed in step 12. The most computationally de-
 305 manding part of condensation is the factorisation of the inner-element matrix
 306 $\underline{\mathbf{A}}_{ii}^{-1}$. It is known [60] that the standard factorisation (including LU decom-
 307 position) algorithms have a cost $\mathcal{O}(n^3)$. Considering that the size of $\underline{\mathbf{A}}_{ii}$ can
 308 be described with equations (7) and (9), the resulting cost of factorising this
 309 matrix is $N_{el}N_{eq}^3(P-1)^9$ in 3D and $N_{el}N_{eq}^3(P-1)^6$ in 2D.

310 The second important operation is the Sparse Matrix-Matrix multiplica-
 311 tions (SpGEMM). In our computations we rely on PETSc libraries to perform
 312 SpGEMM on compressed sparse row matrices. An upper bound for the cost
 313 of for matrix-matrix SpGEMM can be easily calculated assuming n matrix-
 314 vector SpMV. If the sparse matrix has nnz non-zero entries, then the matrix-
 315 matrix cost scales as $\mathcal{O}(n \times nnz)$. This estimate is not optimal and improved
 316 algorithms can be found in the literature [61, 62, 63], but this upper bound is
 317 accurate enough to analyse our condensed costs. To compute the condensed
 318 system, we perform two SpGEMM operations to compute $\underline{\mathbf{A}}_{bi}\underline{\mathbf{A}}_{ii}^{-1}\underline{\mathbf{A}}_{ib}$. We
 319 assume that $\underline{\mathbf{A}}_{ii}^{-1}$ has dense blocks of size $nb_{ii} = N_{eq}(P-1)^d$ and that the num-
 320 ber of non-zeros is larger in $\underline{\mathbf{A}}_{ii}^{-1}$ than in the very sparse $\underline{\mathbf{A}}_{ib}$ (see Appendix
 321 C.27 for the estimation of the number of non-zeros in off-diagonal blocks of
 322 the Jacobian matrix, which scales as $N_{eq}^2(P+1)^2(4P+1)$). Taking into ac-

323 count that the size of the blocks of the Schur complement is $nb_{bb} = N_{eq}(6P^2 +$
 324 $2)$ in 3D and $nb_{bb} = N_{eq}4P$ in 2D, we approximate the cost of the SpGEMM
 325 operation as $\mathcal{O}(N_{el}N_{eq}^3(6P^2 + 2)(P - 1)^6)$ in 3D and $\mathcal{O}(N_{el}N_{eq}^34P(P - 1)^4)$ in
 326 2D. These upper bounds for matrix-matrix SpGEMM show that the inversion
 327 of the matrix $\underline{\mathbf{A}}_{ii}^{-1}$, which scales as $\mathcal{O}((P - 1)^9)$ in 3D and as $\mathcal{O}((P - 1)^6)$ in
 328 2D is the dominant cost in calculating the Schur complement and obtaining
 329 the condensed system.

330 Finally, let us note that the estimation for nnz in Appendix C provides an
 331 upper bound that assumes full coupling between conservative variables. The
 332 real non-zero entries of $\underline{\mathbf{A}}_{bi}$ and $\underline{\mathbf{A}}_{ib}$ have few non-zeros, therefore in practical
 333 computations one would always expect a lower computational costs.

334 3.2. Cost of factorising the preconditioner

335 After computing the condensed system $\underline{\mathbf{A}}_{cond}$ in Algorithm 1 (step 8), we
 336 compute the preconditioner (step 9). As mentioned in section 2.4, we employ
 337 an element Block-Jacobi preconditioner to speed-up the convergence. If the
 338 full system (4) is considered, we factorise the whole element-blocks of matrix
 339 $\underline{\mathbf{A}}$ of size $N_{el}nb$, which has an operation count of $N_{el}N_{eq}^3(P + 1)^9$ in 3D and
 340 $N_{el}N_{eq}^3(P + 1)^6$ in 2D. If the condensed system is considered, we factorise
 341 the skeleton-element blocks of matrix $\underline{\mathbf{A}}_{cond}$ of size $N_{el}nb_{bb}$, which has a cost
 342 $N_{el}N_{eq}^3 [(P + 1)^d - (P - 1)^d]^3$. This can be simplified to $N_{el}N_{eq}^3(6P^2 + 2)^3$
 343 in 3D and $N_{el}N_{eq}^3(4P)^3$ in 2D. The cost of factorising the preconditioner
 344 is henceforth referred to as *preconditioner cost*. At this stage, we can al-

345 ready foresee that the cost of preconditioning the condensed system is much
 346 cheaper, since it scale as $\mathcal{O}(P^6)$ whilst for the full the cost scales as $\mathcal{O}(P^9)$.

347 Pardo et al. [24] concluded that their *hp*-FEM static condensation im-
 348 plementation for single, linear, second order PDE was computationally more
 349 efficient than the full system of equations when the number of iterations is
 350 high enough, since shorter times per iteration compensate the *condensation*
 351 *cost*. For time-dependent problems, like the compressible flow simulations
 352 considered here, this cost becomes even less important, as we can store the
 353 condensed matrix (in our matrix-based approach) and re-use it.

354 3.3. Cost of the preconditioned-GMRES solver

355 Step 13 in Algorithm 1 is detailed in Algorithm 2 where a preconditioned
 356 version of GMRES developed by Saad and Schultz [64] is presented. This is
 357 implemented in the PETSc library [49, 50, 51] and has been used in this work.
 358 In Algorithm 2, \mathbf{R} and \mathbf{V} represent the residual and its normalised version.
 359 m is dimension of the Krylov subspace $\underline{\mathbf{W}}_m$ with orthonormal vectors \mathbf{W}_j
 360 and $\underline{\mathbf{H}}_m$ is the reduced Hessenberg matrix. $\underline{\mathbf{A}}$, $\Delta\mathbf{Q}$ and \mathbf{B} represents either
 361 the full Jacobian matrix $\underline{\mathbf{A}}$, approximate solution $\Delta\mathbf{Q}$ and the right had
 362 side (RHS) \mathbf{B} for the full system. Alternatively, when the condensed system
 363 is solved, we use the condensed Jacobian $\underline{\mathbf{A}}_{cond}$, $\Delta\mathbf{Q}_b$ and condensed RHS
 364 \mathbf{B}_{cond} .

Algorithm 2 Preconditioned GMRES-Solver

```
1: function GMRES-SOLVE( $\Delta\mathbf{Q}$ ,  $\underline{\mathbf{A}}$ ,  $\underline{\mathbf{P}}^{-1}$ ,  $\mathbf{B}$ )
2:    $\mathbf{R}_0 \leftarrow \mathbf{B} - \underline{\mathbf{A}}\Delta\mathbf{Q}$ 
3:    $\mathbf{V}_1 \leftarrow \mathbf{R}_0 / \|\mathbf{R}_0\|_2$ 
4:   for  $j = 1, \dots, m$  do
5:      $\mathbf{Z}_j \leftarrow \underline{\mathbf{P}}^{-1}\mathbf{V}_j$ 
6:      $\mathbf{W} \leftarrow \underline{\mathbf{A}}\mathbf{Z}_j$ 
7:      $\underline{\mathbf{H}}_{i,j} \leftarrow \mathbf{W}^T\mathbf{V}_i$ ,  $i = 1, \dots, j$ 
8:      $\mathbf{W} \leftarrow \mathbf{W} - \sum_{i=1}^j \underline{\mathbf{H}}_{i,j}\mathbf{V}_i$ 
9:      $\mathbf{W} \leftarrow \underline{\mathbf{H}}_{j+1,j} / \|\mathbf{W}\|_2$ 
10:     $\mathbf{V}_{j+1} \leftarrow \mathbf{W} / \underline{\mathbf{H}}_{j+1,j}$ 
11:     $\Delta\mathbf{Q} \leftarrow \Delta\mathbf{Q} + \mathbf{Z}_m\mathbf{Y}_m$ , where  $\mathbf{Y}_m$  minimizes  $\|\beta\mathbf{e}_1 - \underline{\mathbf{H}}_m\mathbf{Y}\|$ 
12: return  $\Delta\mathbf{Q}$ 
```

365 The main costs within the GMRES iterative solver, arise from Sparse
366 Matrix-Vector products (SpMV) (see steps 5 and 6 of Algorithm 2), which
367 are governed by the number of non-zero entries nnz [65], in matrices $\underline{\mathbf{P}}^{-1}$ and
368 $\underline{\mathbf{A}}$ [65]. Note that each nnz performs one multiplication and one addition, and
369 we omit operation counts related to loading/storing variables. In addition to
370 SpMV operations, GMRES also incorporates a large amount of purely vector
371 operations (mainly dot products used to update the Hessenberg matrix, step
372 in Algorithm 2). Their cost is proportional to the matrix size n , and have
373 typically lower cost than sparse matrix-vector products. Therefore we focus
374 only on SpMV operations.

375 The cost of Jacobian-SpMV (Step 6) is a function of nnz_{full} for full system
376 $\underline{\mathbf{A}}$ and nnz_{cond} for the condensed system $\underline{\mathbf{A}}_{cond}$. In Appendix C, we have
377 detailed the derivation of an upper bound for the number of non-zero entries
378 for the Jacobian DGSEM matrix, see table 1 and Appendix C. Similarly, we

379 also express the number of non-zero entries in the condensed matrix nnz_{cond} ,
 380 see equation (C.30) in Appendix C. This enables the calculation the costs of
 381 step 6: Precondition-SpMV $\mathbf{Z} = \underline{\mathbf{P}}^{-1}\mathbf{V}$ and Jacobian-SpMV $\mathbf{W} = \underline{\mathbf{A}}\mathbf{Z}$ in
 382 terms of (P, N_{eq}, d) , as summarised in Table 2. Since the preconditioner is a
 383 locally dense matrix (block diagonal part is dense, while the off-diagonal parts
 384 are empty), we can bound the number of non-zero entries by the number of
 385 total entries in the diagonal blocks $nnz = N_{el}nb^2$. Therefore, the cost of the
 386 preconditioner-SpMV $\mathbf{Z} = \underline{\mathbf{P}}^{-1}\mathbf{V}$, presented in step 5 in Algorithm 2 can be
 387 expressed as $N_{el}N_{eq}^2(P+1)^6$ in 3D and $N_{el}N_{eq}^2(P+1)^4$ in 2D, if the full system
 388 is considered. For the condensed system, the costs are $N_{el}N_{eq}^2(6P^2 + 2)^2$ and
 389 $N_{el}N_{eq}^2 16P^2$ for 3D and 2D, respectively. The main preconditioned-GMRES
 390 costs are included in Table 2.

391 These estimates show asymptotic advantages for the condensed system,
 392 as P increases, for the two main steps within the preconditioned GMRES
 393 solver, further discussion can be found in the next section. In Section 4,
 394 Figures 2a and 6a report measured computational costs of GMRES in detail
 395 for the range of polynomial orders $P = 2, \dots, 8$. The cost of GMRES (step
 396 13 in Algorithm 1) is referred to as *solver cost*, in the following sections.

397 3.4. Summary of computational costs

398 Table 2 presents a summary of the estimated costs for the essential oper-
 399 ations considered in the time stepping algorithm Algorithm 1, including the
 400 preconditioned-GMRES main steps. The biggest computational effort relates

401 to the factorisation of element-blocks needed to factorise the preconditioner
402 for the full system and inner-element matrix $\underline{\mathbf{A}}_{ii}$ for the Schur complement,
403 both scaling as $\mathcal{O}(P^9)$. As shown in the Table 2, factorising the blocks for the
404 condensed preconditioner has a significant lower cost $\mathcal{O}(P^6)$. Similarly, the
405 main GMRES steps favor from the use of static condensation. In 3D, both
406 steps scale as $\mathcal{O}(P^4)$ for the condensed system, whilst they scale as $\mathcal{O}(P^5)$
407 and $\mathcal{O}(P^6)$ for the full system. These advantages are also expected in 2D
408 simulations.

Table 2: Summary of the estimated leading costs of main operations in Algorithm 1 for 2D and 3D. Full and condensed systems are included.

3D		
	Full system	Condensed system
$\underline{\mathbf{A}}_{ii}^{-1}$	-	$N_{el}N_{eq}^3(P-1)^9$
SpGEMM	-	$N_{el}N_{eq}^3P^8$
$\underline{\mathbf{P}}^{-1}$	$N_{el}N_{eq}^3(P+1)^9$	$N_{el}N_{eq}^3P^6$
GMRES $\underline{\mathbf{A}}\mathbf{z}$	$N_{el}N_{eq}^2P^5$	$25N_{el}N_{eq}^2P^4$
GMRES $\underline{\mathbf{P}}^{-1}\mathbf{v}$	$N_{el}N_{eq}^2(P+1)^6$	$N_{el}N_{eq}^26P^4$
2D		
	Full system	Condensed system
$\underline{\mathbf{A}}_{ii}^{-1}$	-	$N_{el}N_{eq}^3(P-1)^6$
SpGEMM	-	$N_{el}N_{eq}^3P^5$
$\underline{\mathbf{P}}^{-1}$	$N_{el}N_{eq}^3(P+1)^6$	$N_{el}N_{eq}^364P^3$
GMRES $\underline{\mathbf{A}}\mathbf{z}$	$N_{el}N_{eq}^2P^4$	$13N_{el}N_{eq}^2P^2$
GMRES $\underline{\mathbf{P}}^{-1}\mathbf{v}$	$N_{el}N_{eq}^2(P+1)^4$	$N_{el}N_{eq}^216P^2$

409 In Section 4, we study the difference in computational costs for both,
410 full and condensed Block-Jacobi preconditioners. There simulated costs are
411 compared to the summarised estimated. We present the results in Figures 3a
412 and 7a together with the *condensation costs*. Finally, we note that the use of

413 block preconditioners, that exploit the structure of DGSEM, has proven to be
 414 an important part in obtaining faster convergence rates for DG based solvers
 415 [33, 31, 32, 35, 36]. It has been advocated that Block-Jacobi preconditioner
 416 do not scale well in DG, which is indeed the case for the full system, since
 417 the block size scales with $(P + 1)^3$, and associated cost $\mathcal{O}(P^9)$. However, the
 418 static condensed block size scales with $6P^2 + 2$ with costs $\mathcal{O}(P^6)$ in 3D and
 419 with $4P$ and cost $\mathcal{O}(P^3)$ in 2D, which renders Block-Jacobi preconditioner an
 420 interesting scalable preconditioner for the condensed GL-DGSEM approach.

421 4. Numerical results

422 We consider two test cases: a 3D manufactured solution and a 2D flow
 423 over NACA0012 airfoil at a high Angle of Attack (AOA) leading to an un-
 424 steady regime. The manufactured solution case illustrates the use of implicit
 425 time-marching solvers to reach a steady state solution, whilst the NACA0012
 426 test case quantifies the improved cost in an unsteady flow simulation, with
 427 vortex shedding. The Mach number is set to $\text{Ma}=0.1$ for manufactured solu-
 428 tion problem (other Ma and Re can be found in Appendix A) and $\text{Ma}=0.3$
 429 for the NACA cases. For all the steady cases, we fix the final residual of the
 430 simulations to $\|\underline{\mathbf{M}}^{-1}\mathbf{F}(\mathbf{Q}) - \mathbf{S}\|_\infty = 10^{-8}$ (see Algorithm 1) such that we
 431 compare the various schemes for the same accuracy.

432 The objective of the test cases is to validate the theoretical findings pre-
 433 sented in the previous section. Therefore the main costs of the time marching
 434 scheme (see Algorithm 1) are compared with the theoretical cost estimations

435 (summarised in Table 2) for the two test cases. Additionally, the total cost
 436 to perform the simulations is included, to quantify the overall efficiency of
 437 the implicit statically condensed system compared to the full system.

438 4.1. Steady simulation: Manufactured Solution

439 The manufactured solution case is obtained by selecting an exact solution
 440 to the compressible Navier-Stokes equations:

$$\begin{aligned} \rho &= p = e^{-5 \cdot (4(x-\frac{1}{2})^2 + (y-\frac{1}{2})^2 + (z-\frac{1}{2})^2)} + 1, \\ u &= v = w = 1, \end{aligned} \tag{11}$$

441 to then extract the balancing source terms:

$$\mathbf{s} = \begin{bmatrix} s_\rho \\ s_{\rho u} \\ s_{\rho v} \\ s_{\rho w} \\ s_{\rho E} \end{bmatrix} = \begin{bmatrix} 40(x - \frac{1}{2}) + 10(y - \frac{1}{2}) + 10(z - \frac{1}{2}) \\ 80(x - \frac{1}{2}) + 10(y - \frac{1}{2}) + 10(z - \frac{1}{2}) \\ 40(x - \frac{1}{2}) + 20(y - \frac{1}{2}) + 10(z - \frac{1}{2}) \\ 40(x - \frac{1}{2}) + 10(y - \frac{1}{2}) + 20(z - \frac{1}{2}) \\ [40(x - \frac{1}{2}) + 10(y - \frac{1}{2}) + 10(z - \frac{1}{2})] [\frac{5}{2} + \frac{1}{\gamma-1}] \end{bmatrix} \cdot e^{-5(4(x-\frac{1}{2})^2 + (y-\frac{1}{2})^2 + (z-\frac{1}{2})^2)}. \tag{12}$$

442 We select the computational domain to be a $[0, 1]^3$ cube with 64 hexa-
 443 hedral uniform elements. The solutions to the compressible Navier-Stokes
 444 equations (1) with the source term (12) can be seen in the Figure 1. Neither

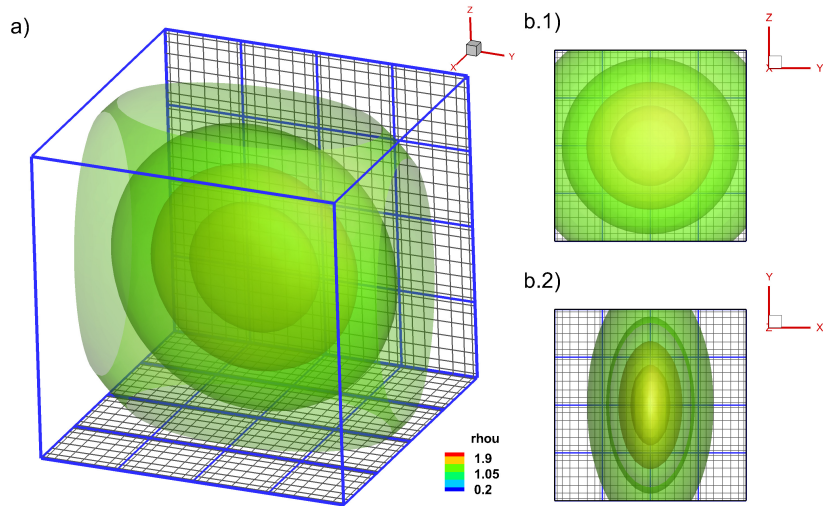


Figure 1: 3D Manufactured Solution: Solution of x-momentum ρu for a mesh of 64 hexahedral and polynomial order $P=8$. Figure a) 3D view, Figure b.1) and b.2) show cross-sections of yz and xy planes, respectively.

445 Mach nor Reynolds numbers have an impact on the final solution, however,
 446 both those parameters have a strong influence on the numerical scheme. We
 447 set the Reynolds number to $Re = 1000$ and the Mach to $Ma = 0.1$, but
 448 results for other Re and Ma can be found in Appendix B. The time-step
 449 size in the implicit computations is set $\Delta t = 0.1$ for all the polynomial orders
 450 and both systems.

451 Table 3 presents a summary of the conducted simulations for the full and
 452 condensed systems. We include the number of Jacobian updates i_{Jac} (iden-
 453 tical for both full and condensed systems), the averaged number of Newton
 454 iterations per one time step $\frac{i_{Newton}}{i_{\Delta t}}$, the averaged number of GMRES itera-
 455 tions per one Newton solve $\frac{i_{GMRES}}{i_{Newton}}$ along with number of non-zero entries in
 456 full nnz_{full} and condensed nnz_{cond} systems. We observe that the number of

457 Newton iterations per one time step $\frac{i_{Newton}}{i_{\Delta t}}$ is constant for all polynomials
458 and almost identical for the full and condensed system, consequently with us-
459 ing the same number of Jacobian updates in the full and condensed systems.
460 The averaged number of GMRES iterations per one Newton solve $\frac{i_{GMRES}}{i_{Newton}}$,
461 increases when using higher polynomial orders, scales similarly for both full
462 and condensed systems. We also observe that the number on non-zeros is
463 larger for the condensed system. This is not the expected behaviour for high
464 polynomials, but due to the tight coupled stencil of the condensed system,
465 this can be expected for low polynomial orders.

Table 3: 3D Manufactured Solution: Number of Jacobian updates i_{Jac} (identical for both full and condensed systems), averaged number of Newton iterations per one time step $\frac{i_{Newton}}{i_{\Delta t}}$ and averaged number of GMRES iterations per one Newton solve $\frac{i_{GMRES}}{i_{Newton}}$ along with number of non-zero entries in full nnz_{full} and condensed nnz_{cond} systems. For all cases considered in the table number of time steps needed to reach the steady state is $i_{\Delta t} = 50$, for polynomial orders $P = 2, \dots, 8$.

P	i_{Jac}	Full system		Condensed system		Nonzero entries	
		$\frac{i_{Newton}}{i_{\Delta t}}$	$\frac{i_{GMRES}}{i_{Newton}}$	$\frac{i_{Newton}}{i_{\Delta t}}$	$\frac{i_{GMRES}}{i_{Newton}}$	nnz_{full}	nnz_{cond}
2	3	6.4	3.5	6.4	3.5	6.5×10^5	1.2×10^6
3	3	6.1	4.4	6.1	4.4	2.2×10^6	6.6×10^6
4	3	5.9	5.7	6.1	5.5	6.0×10^6	2.2×10^7
5	3	6.1	6.7	6.1	6.5	1.3×10^7	5.6×10^7
6	4	6.3	7.7	6.4	7.6	2.8×10^7	1.2×10^8
7	4	6.5	8.6	6.6	8.5	5.3×10^7	2.2×10^8
8	5	6.6	10.0	6.6	9.9	9.4×10^7	3.9×10^8

466 Although the averaged linear solver iteration count is the same for both
467 systems, this can be interpreted as an advantage of using static condensation
468 with the cheaper skeleton-element Block-Jacobi. The similar iteration count
469 has been observed in the past [24, 66] for finite element formulations (and

470 moderate polynomials P). There, the authors argued that even if the condi-
 471 tion number of condensed Jacobian scales much better with P , the spectral
 472 radius of the iteration matrix, with a good preconditioner, is very similar for
 473 both systems, leading to similar number of iterations. Coherently with the
 474 findings of the aforementioned publications, we find almost the same number
 475 of iterations for full and condensed systems, but the latter being cheaper due
 476 to its smaller size, see Figure 2b.

477 The table is completed with Figures 2 and 3, where the total GMRES cost,
 478 the averaged solver cost per one linear system solve, the timing of factorising
 479 the preconditioner and the total simulation cost are depicted for the full
 480 and condensed systems and for polynomial orders $P = 2, \dots, 8$. The figures
 481 include the slopes for the theoretical estimates found in previous sections.
 482 Figure 2a splits the solver costs into the two main preconditioned-GMRES
 483 solver steps: preconditioner-SpMV $T_{\underline{P}^{-1}\mathbf{v}}$ and Jacobian-SpMV $T_{\underline{A}\mathbf{z}}$. Note
 484 that the rest of the GMRES costs are negligible. As estimated in Section 3,
 485 $T_{\underline{A}\mathbf{z}}$ is larger for the condensed system due to higher number of non-zeros nnz ,
 486 however the preconditioner-SpMV $T_{\underline{P}^{-1}\mathbf{v}}$ is much cheaper and compensates
 487 $T_{\underline{A}\mathbf{z}}$, which results in faster overall iterations. Additionally, the advantage of
 488 using static condensation in terms of solver costs becomes more noticeable
 489 for high polynomial orders. In all cases, the theoretical estimates are in good
 490 agreement with the numerical results.

491 Figure 3 presents the factorisation costs of the preconditioner along with
 492 condensation cost and the total time of the simulation. The factorisation of

493 the preconditioner matches well the theoretical estimates (see Table 2) for
 494 high enough polynomial orders. Discrepancies at low orders are attributed
 495 to the relatively small 3D problem considered and the effect of boundary
 496 conditions. In any case, it can be seen that despite the cost of condensing the
 497 system, the solver cost benefits from the condensation (Figure 2b), leading to
 498 overall faster solves, which illustrates the beneficial effect of using a condensed
 499 system for the higher polynomial orders.

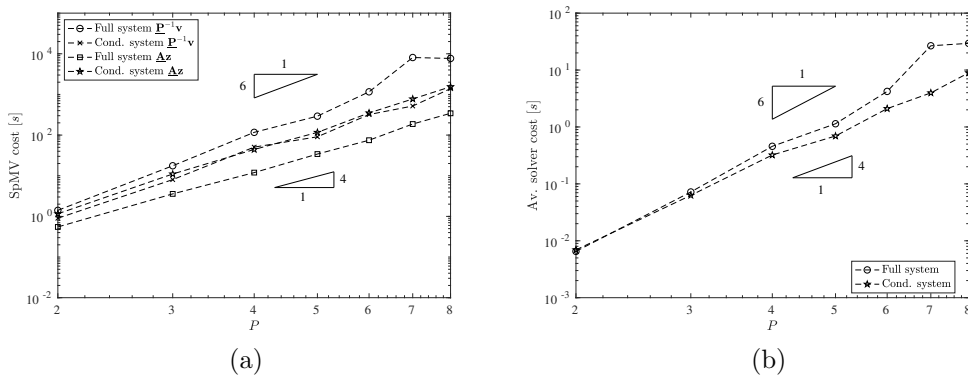


Figure 2: 3D Manufactured Solution: a) Total cost of the GMRES split in two major operations (in seconds) and b) Averaged GMRES solver cost (in seconds) per one linear system solve, for full and condensed systems for polynomial order $P = 2, \dots, 8$. Theoretical slopes are included depicted with a triangle.

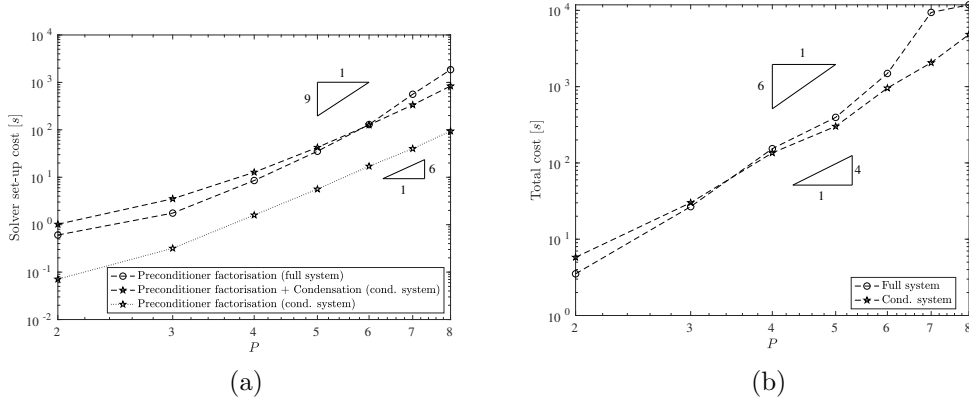


Figure 3: 3D Manufactured Solution: a) Timing of factorising the preconditioner (in seconds) and b) Total simulation cost (in seconds) to reach a tolerance 1×10^{-8} for the full and condensed systems for $P = 2, \dots, 8$. Theoretical slopes are included depicted with a triangle.

500 For completeness, we include a brief study for this problem, in Ap-
 501 pendix B, for a range of Mach numbers, $0.1 \leq \text{Ma} \leq 0.8$ and Reynolds numbers,
 502 $200 \leq \text{Re} \leq 1000$ and show that the advantages of the static condensation are
 503 maintained for a wide range of flow conditions, and for a variety of polynomial
 504 orders.

505 4.2. Unsteady simulation: NACA0012 at $AOA = 20^\circ$

506 In this section, we challenge the static condensation technique for un-
 507 steady flows with application to aerodynamics. We simulate an unsteady
 508 NACA0012 case using a 2D computational squared domain of size 20×20
 509 chords, with 1730 quadrilateral elements. Figure 4 depicts the h -mesh (in
 510 black) and the Gauss-Lobatto mesh (in gray) near the NACA0012 airfoil,
 511 and also the contours of x-momentum for the wake flow. To trigger vortex

512 shedding and study the performance of the implicit time-marching method
513 for unsteady regimes, we set the Reynolds number to $Re=200$ and the angle
514 of attack to $AOA = 20^\circ$ (see Figure 4).

515 In steady problems, one of the main advantages of implicit time-integration
516 schemes is that it is possible to increase the time-step size several orders of
517 magnitude without losing accuracy or affecting stability [67]. However, in
518 unsteady simulations the time-step size is bounded by accuracy constraints.
519 This means that the time step in the implicit time-marching schemes has
520 to be low enough to capture the physics of the problem, hence the perfor-
521 mance of implicit time-marching schemes depends on the underlying physical
522 problem at hand. In NACA0012, the characteristic physical time (one vortex
523 shedding cycle) is 200 times larger than the time step selected for the implicit
524 time-marching scheme. This restriction precludes the use of very large time
525 steps in implicit solvers. For this reason, in the unsteady case, the implicit
526 time step as been restricted to maintain accuracy (as shown in Figure 5). The
527 time-step size in the explicit computations (ERK3), provided as reference for
528 accuracy, is limited to $\Delta t = 2.0 \times 10^{-5}$, which is the maximum permitted by
529 stability constraints for $P = 5$. In contrast, the time-step size in the implicit
530 computations is set to $\Delta t = 1.0 \times 10^{-2}$, which is sufficiently low to capture
531 the flow features accurately.

532 In this section, we show that the statically condensed DGSEM is able to
533 outperform the standard full system for the same step size and that both
534 methods provide accurate results. We provide results using an explicit RK3




Figure 4: Unsteady 2D NACA0012: unsteady flow at $Re=200$ and $AOA = 20^\circ$. Zoomed regions showing h -mesh (in black) and Gauss-Lobatto mesh (in gray) in a) and wake flow field in b). All figures include x-momentum contours.

535 scheme as a reference. The comparison shows that the additional operations
 536 necessary to calculate the Schur complement, in the condensed system, do
 537 not damage the accuracy of the final solution with round-off errors.


538 Before proceeding with the costs, we present comparisons for the schemes
 539 in terms of accuracy, in Table 4. We simulate the unsteady flow for 10 vortex
 540 shedding cycles and compute mean lift, mean drag, and the Strouhal number.
 541 Let us note that once the polynomial order is fixed, the differences in mean
 542 lift, mean drag and Strouhal are negligible (i.e. below 10^{-5}) when using
 543 different time-marching schemes.

544 For completeness, Figure 5 depicts drag and lift curves for $P = 3$, com-
 545 puted with the explicit and implicit methods. We observe that explicit and
 546 implicit results match remarkably well, illustrating that there is no loss of
 547 accuracy when using implicit time-marching with moderate time steps.

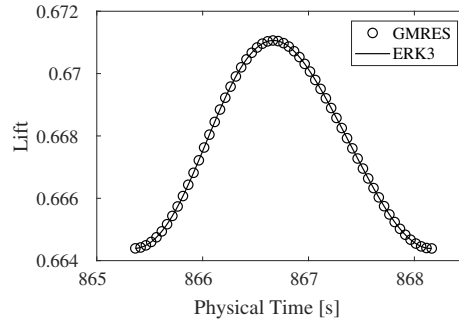
Table 4: Unsteady 2D NACA0012: Time step, mean drag, mean lift and Strouhal number St ; for explicit (ERK3) and implicit solver (GMRES) and polynomial orders $P = 2, 3, 4, 5$.

	$P = 2$			$P = 3$		
	ERK3	BDF2 full	BDF2 <i>cond.</i>	ERK3	BDF2 full	BDF2 <i>cond.</i>
Δt	2.7×10^{-5}	1.0×10^{-2}	1.0×10^{-2}	2.7×10^{-5}	1.0×10^{-2}	1.0×10^{-2}
Mean Drag	0.4383	0.4383	0.4383	0.4342	0.4342	0.4342
Mean Lift	0.6753	0.6753	0.6753	0.6677	0.6677	0.6677
St	0.3535	0.3530	0.3530	0.3565	0.3565	0.3565

	$P = 4$			$P = 5$		
	ERK3	BDF2 full	BDF2 <i>cond.</i>	ERK3	BDF2 full	BDF2 <i>cond.</i>
Δt	2.7×10^{-5}	1.0×10^{-2}	1.0×10^{-2}	2.7×10^{-5}	1.0×10^{-2}	1.0×10^{-2}
Mean Drag	0.4345	0.4345	0.4345	0.4342	0.4342	0.4342
Mean Lift	0.6664	0.6664	0.6664	0.6651	0.6651	0.6651
St	0.3577	0.3576	0.3576	0.3558	0.3558	0.3558



(a)



(b)

Figure 5: Unsteady 2D NACA0012: Close-up comparison of explicit and implicit results for drag and lift for a single shedding cycle.

548 We now explore the different costs. Table 5 shows detailed information
 549 about number of Jacobian updates i_{Jac} , the averaged number of Newton
 550 iterations per one time step $\frac{i_{Newton}}{i_{\Delta t}}$ and averaged number of GMRES itera-
 551 tions per one Newton solve $\frac{i_{GMRES}}{i_{Newton}}$. As in the previous Manufactured Solu-
 552 tion problem, conducting the simulation based on a smaller (but with more
 553 non-zeros) Jacobian matrix has almost no impact in the number of Newton
 554 iterations. Also like in the previous steady-state case, the averaged num-
 555 ber of GMRES iterations is similar for both systems, but the iterations are,
 556 again, more efficient for the condensed system (Figure 6b). Unlike in the
 557 previous problem, the solver set-up costs (factorisation and condensation)
 558 do not constitute a big portion of the total simulation time, see Figure 7a,
 559 thus the advantage for the condensed system is clearly seen in Figure 7b.
 560 This is due to the fact that the Jacobian matrix is updated less frequently in
 561 this problem, and therefore the relative cost of the solver set-up in the total
 562 simulation cost is smaller. For this particular test case and range of poly-
 563 nomial orders, the solver set-up cost for the full system is cheaper than the
 564 theoretical prediction. However, it is still more costly than the condensation
 565 cost.

566 It can be seen that the static-condensation method provides the same
 567 accuracy up to given tolerance as the full system, but it is up to 40 % faster
 568 for the highest polynomial orders ($P = 4, 5$). As in the previous section,
 569 we also present the detailed results of the solver cost, Figures 6a and 6b.
 570 Again, the condensed system has more non-zeros nnz (Table 5), but the faster

571 preconditioner-SpMV compensates this cost and leads to faster simulations.
 572 Theoretical and measured preconditioner-SpMV operations for both systems
 573 agree well.

574 Finally, we can conclude that our static condensation time-marching
 575 method is more efficient for large polynomials, than the full system tech-
 576 nique, even for unsteady problems, whilst providing accurate results.

Table 5: Unsteady 2D NACA0012: Number of Jacobian updates i_{Jac} (computed only once and identical for both full and condensed systems), averaged number of Newton iterations per one time step $\frac{i_{Newton}}{i_{\Delta t}}$ and averaged number of GMRES iterations per one Newton solve $\frac{i_{GMRES}}{i_{Newton}}$ along with number of non-zero entries in full nnz_{full} and condensed nnz_{cond} systems. For all cases considered in the table number of time steps needed to compute one cycle is $i_{\Delta t} = 280$, for polynomial orders $P = 2, \dots, 5$.

P	i_{Jac}	Full system		Condensed system		Nonzero entries	
		$\frac{i_{Newton}}{i_{\Delta t}}$	$\frac{i_{GMRES}}{i_{Newton}}$	$\frac{i_{Newton}}{i_{\Delta t}}$	$\frac{i_{GMRES}}{i_{Newton}}$	nnz_{full}	nnz_{cond}
2	1	5.5	5.0	5.5	5.0	1.9×10^7	4.9×10^7
3		11.2	10.0	11.1	9.8	4.5×10^7	1.2×10^8
4		11.7	11.7	11.7	11.5	9.0×10^7	2.4×10^8
5		11.6	13.5	11.6	13.2	1.5×10^8	3.9×10^8




Figure 6: Unsteady 2D NACA0012: a) Total cost of the GMRES split in two major operations (in seconds) and b) Averaged GMRES solver cost (in seconds) per one linear system solve, for the full and condensed systems for $P = 2, \dots, 5$. Theoretical slopes are included depicted with a triangle.




Figure 7: Unsteady 2D NACA0012: a) Timing of factorising the preconditioner and b) Total simulation cost (in seconds) to complete one shedding period for full and condensed systems for $P = 2, \dots, 5$. Theoretical slopes are included depicted with a triangle.

577 **5. Conclusion**

578 We have analysed the advantages of performing static condensation on the
 579 compressible Navier-Stokes equations discretised with DGSEM and Gauss-

580 Lobatto points. The work focuses on the implicit preconditioned-GMRES
 581 time discretisations, and we have compared computational costs of solving
 582 the standard full Jacobian system to the static condensation technique for
 583 GL-DGSEM, detailed in Rueda-Ramírez et al. [1], both preconditioned us-
 584 ing Block-Jacobi. To allow for fair comparisons, we split the costs into three
 585 categories: computation of the preconditioner, condensation costs for the
 586 statically condensed system and the solver GMRES cost to solve the full and
 587 condensed systems. We compare our numerical results with theoretical com-
 588 putational costs (Table 2), which include unpublished estimates for DGSEM.
 589 The theoretical estimates agree well with our simulations and provide solid
 590 bases for understanding the different costs involved.

591 For all cases included (steady-state 3D Manufactured Solution and un-
 592 steady 2D NACA0012), the static condensation shows accelerations (for large
 593 polynomial orders) due to the significantly faster solver time per single linear
 594 system solve. The accelerations are up to 30% for the Manufactured Solution
 595 and up to 40% for NACA0012 case, for the highest polynomial considered.
 596 Block-Jacobi preconditioner do not scale well with the polynomial order,
 597 which is indeed the case for the full system, since the element block Jacobian
 598 scales with $(P + 1)^3$. However, we have shown that the statically condensed
 599 block size scales with $6P^2 + 2$ in 3D and with $4P$ in 2D, which renders
 600 Block-Jacobi preconditioner an interesting preconditioner for the condensed
 601 GL-DGSEM approach. Let us note that recent sum-factorisation techniques
 602 have been developed for high polynomials in discontinuous [48] and continu-

603 ous Galerkin [68] approaches that decrease cost of factorising the blocks and
604 show improved scalings for Block-Jacobi preconditioners. In the future, this
605 approach may be applied to decrease the computational cost of condensed
606 systems to further enhance the presented methodology. One drawback as-
607 sociated to the statically condensed system is the additional cost related to
608 assembling the Schur complement (see Section 2.2 for more details). How-
609 ever, this cost is not high enough to mask the advantages of using static
610 condensation, for high polynomial orders.

611 This manuscript compares iterative time-marching methods in serial, to
612 avoid discrepancies due to parallelisation when comparing implicit tech-
613 niques. Taking into account that Block-Jacobi preconditioners can be eas-
614 ily parallelised, we expect that future parallel implementation will lead to
615 cheaper parallelised costs and less communication than when using the full
616 system, as well as lower memory requirements. Future work, will assess
617 the improvements in performance of implicit schemes (and especially of the
618 static condensation methods) in many-core parallel environments and with
619 more sophisticated preconditioners, including multilevel p -multigrid, specifi-
620 cally tailored for statically condensed systems.

621 **Acknowledgements**

622 Wojciech Laskowski and Esteban Ferrer would like to thank the Euro-
623 pean Union's Horizon 2020 Research and Innovation Program under the
624 Marie Skłodowska-Curie grant agreement No 813605 for the ASIMIA ITN-

625 EID project. Additionally, Andrés Rueda-Ramírez acknowledges the fund-
626 ing received by the project SSeMID under the Marie Skłodowska-Curie grant
627 agreement No 675008. Gonzalo Rubio acknowledges the funding received
628 by the grant *Ayudas dirigidas al PDI para el fomento de la participación*
629 *en solicitudes de proyectos H2020* from Universidad Politécnica de Madrid.
630 Finally, the authors gratefully acknowledge the Universidad Politécnica de
631 Madrid (www.upm.es) for providing computing resources on Magerit Super-
632 computer. Finally, the authors thank the reviewers for suggesting a deeper
633 analysis by means of theoretical estimates, which have clearly improved the
634 manuscript.

635 **Appendix A. Preliminary assessment of preconditioners**

636 In this section, we perform a preliminary study to assess the efficiency
637 of several preconditioners: Block Jacobi and Incomplete LU (i.e. ILU(0),
638 ILU(1) and ILU(2)). We compare the effect of the preconditioning and re-
639 ordering in both the full and statically condensed systems.

640 For this preliminary selection, a manufactured with 8 hexahedral elements
641 is selected. This case is smaller than the one considered in Section 4.1. We
642 also increased the tolerance for the linear solver to $TOL_{GMRES} = e \cdot 0.3^i$
643 along with decreasing time step to $dt = 1e - 2$ for more accurate results. The
644 source term and rest of the parameters are maintained and can be found in
645 Section 4.1.

646 Figure A.8.a, Figure A.8.b and Figure A.8.c show the average number of

647 iterations per Newton-Raphson step, the average solver cost and the cost of
648 factorising the preconditioner. As can be seen in Figure A.8.a, even the sim-
649 plest preconditioners considered (Block-Jacobi and ILU(0)) keep the average
650 number of iterations low, even for high polynomial orders. The number of
651 iterations remains unaltered by the use of static condensation. Figures A.8.b
652 and A.8.c show the averaged solver cost and factorisation cost. Both of them
653 present shorter times for the condensed system than for the full system. The
654 difference between full and condensed systems, in the cost of factorisation for
655 ILU(k), increases when increasing the filling k , as expected for more evolved
656 preconditioners, but we note that the cost is lower for the condensed system,
657 since the system size is smaller. Also the difference between the full and the
658 condensed system in the average solver cost increases for ILU(k) for higher
659 fillings k . Again, the condensed system cost is smaller.

660 From this preliminary analysis, we have chosen the Block-Jacobi for the
661 rest of the paper. The reason is that, although non being optimal in terms
662 of average iteration count, it presents a low memory cost, takes advantage
663 of the element structure in DGSEM and can be easily parallelised, therefore,
664 the results with element Block-Jacobi may provide better bases for further
665 research.

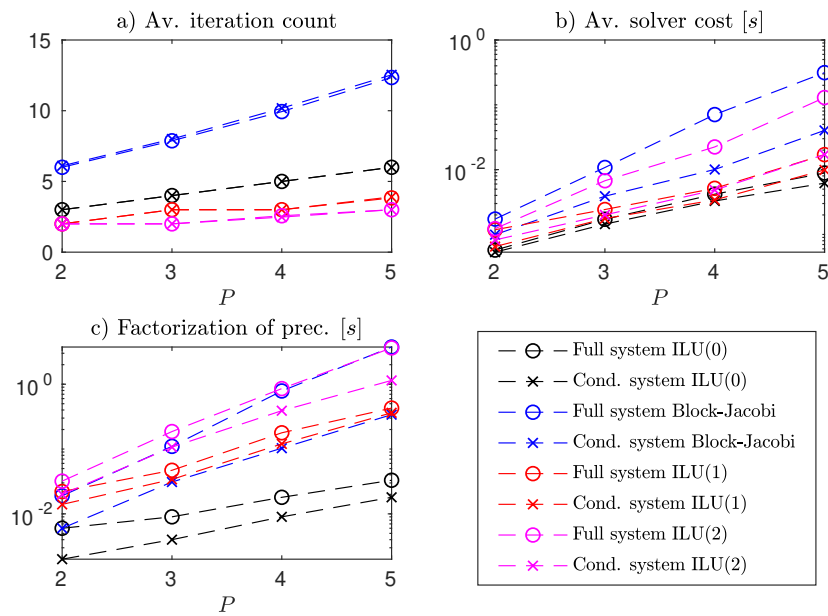


Figure A.8: a) Averaged iteration count (per linear system solve), b) averaged solver cost (per linear system solve) and c) factorisation cost of various types of preconditioners: Block-Jacobi, ILU(0), ILU(1), ILU(2).

666 **Appendix B. Influence of Mach and Reynolds**

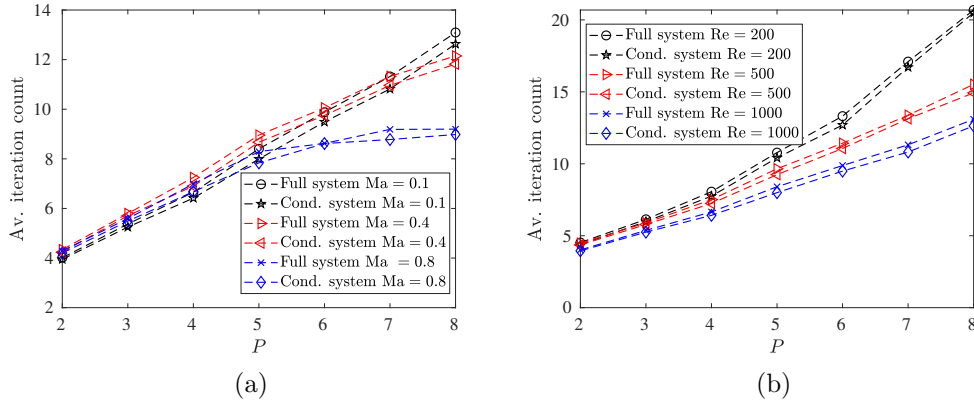


Figure B.9: Influence of a) Mach number (Ma) and b) Reynolds number (Re) on averaged number of iterations. All the different cases for Mach number were simulated with Re = 1000. The Reynolds study was conducted using Ma = 0.1.

667 In this section, we compare full and statically condensed systems for a
 668 range of Mach numbers, $0.1 \leq Ma \leq 0.8$ and Reynolds numbers, $200 \leq Re \leq 1000$.
 669 We use Block-Jacobi preconditioning for all cases.

670 Figure B.9 depicts iteration count for three different Mach numbers B.9a
 671 and three different Reynolds numbers B.9b. As expected, an increase in
 672 the Mach number (in the subsonic range) or in the Reynolds number, has
 673 a positive impact on the averaged iteration count for both condensed and
 674 full systems. We have not observed significant differences in computational
 675 efficiencies. In most cases, the static condensation provided very similar
 676 speed-up as depicted in the Figures 2 and 3. Overall, the static condensation
 677 system shows small improvements, over the full system, for this test case.

678 **Appendix C. Estimation of non-zero entries in the Jacobian Ma-**
679 **trix**

680 The number of non-zero entries in the Jacobian matrix of a DGSEM
681 discretisation depends on the nodes chosen (Gauss or Gauss-Lobatto), on the
682 specifics of the flux (whether it has advection and/or diffusion terms), and
683 on the surface numerical fluxes used. In this section, we derive the analytical
684 expressions for the number of non-zero entries in the Gauss-Lobatto DGSEM
685 Jacobian matrix for systems of nonlinear conservation laws with advection
686 and diffusion terms, and surface numerical fluxes with compact support, as
687 these are the subjects of the present study.

688 To facilitate the analysis, we will note the number of non-zero entries in a
689 single diagonal block of the Jacobian matrix as nnz_d , and the number of non-
690 zero entries in a single off-diagonal block as nnz_o . These expressions should
691 be considered as an upper bound, since the number of non-zeros might be
692 smaller due to the nonlinearities of the problem or the specific properties of
693 the curvilinear mapping, as will be evident in next sections.

A system of nonlinear conservation laws reads

$$\partial_t \mathbf{q} + \vec{\nabla} \cdot (\vec{\mathbf{f}}^a - \vec{\mathbf{f}}^\nu) = \mathbf{0}, \quad (\text{C.1})$$

694 where \mathbf{q} is the state vector of conserved quantities, $\vec{\mathbf{f}}^a$ is the advective flux,
695 and $\vec{\mathbf{f}}^\nu$ is the diffusive flux. Let us analyse the advection and diffusion terms
696 separately.

Given a DGSEM discretisation of the advection terms in (C.1), an entry in the diagonal block that connects degrees freedom h and w of a specific element reads [44, 1]

$$\text{DT}_{hw}^a = - \int_{\Omega} (\underline{\mathbf{J}}^a \phi)_w \cdot \vec{\nabla} \phi_h \, d\Omega + \oint_{\partial\Omega} \hat{\mathbf{f}}_{\mathbf{q}^+}^a \phi_w \phi_h \, dS + \oint_{\partial\Omega \cap \Gamma} \hat{\mathbf{f}}_{\mathbf{q}^-}^a \mathbf{q}_{\mathbf{q}^+}^- \phi_w \phi_h \, dS, \quad (\text{C.2})$$

where Ω is the domain of the element where the degrees of freedom h and w are located, $\partial\Omega$ is the boundary of that domain, $\partial\Omega \cap \Gamma$ is the part of that boundary that belongs to a physical boundary, $\underline{\mathbf{J}}^a = \partial \hat{\mathbf{f}}^a / \partial \mathbf{q}$ is the Jacobian of the advective flux, ϕ_w and ϕ_h are the basis functions that correspond to the degrees of freedom w and h , \mathbf{q}^+ and \mathbf{q}^- are the inner and outer solutions on an element boundary, respectively, $\hat{\mathbf{f}}^a$ is the so-called (advective) surface numerical flux, $\hat{\mathbf{f}}_{\mathbf{q}^\pm}^a$ is its Jacobian with respect to the solution on the element boundary, and $\mathbf{q}_{\mathbf{q}^+}^-$ is the Jacobian of the boundary condition.

The first term of (C.2) generates the densest sparsity. This term can be rewritten using the contravariant fluxes [69] as

$$(\underline{\mathbf{J}}^a \phi)_w \cdot \vec{\nabla} \phi_h = \underbrace{(\tilde{\mathbf{J}}^a \phi)_w \cdot \vec{\nabla}_\xi \phi_h}_{\text{Contravariant form}} = (\tilde{\mathbf{J}}_1^a \phi)_w \frac{\partial \phi_h}{\partial \xi} + (\tilde{\mathbf{J}}_2^a \phi)_w \frac{\partial \phi_h}{\partial \eta} + (\tilde{\mathbf{J}}_3^a \phi)_w \frac{\partial \phi_h}{\partial \zeta}, \quad (\text{C.3})$$

where $\vec{\xi} = (\xi, \eta, \zeta)$ are the coordinates on a reference element $\Omega_\xi = [-1, 1]^3$

that is mapped to physical space with high order polynomials

$$\Omega_\xi \xrightarrow{\vec{x}(\vec{\xi})} \Omega. \quad (\text{C.4})$$

The degrees of freedom indexes h and w can be replaced by the tensor product coordinate indexes $h \leftarrow (i, j, k)$ and $w \leftarrow (r, s, t)$. This allows us to rewrite the basis functions as a tensor product combination of Lagrange interpolating polynomials,

$$\phi_h(\vec{x}(\vec{\xi})) = \ell_i^\xi(\xi) \ell_j^\eta(\eta) \ell_k^\zeta(\zeta) \quad (\text{C.5})$$

$$\phi_w(\vec{x}(\vec{\xi})) = \ell_r^\xi(\xi) \ell_s^\eta(\eta) \ell_t^\zeta(\zeta). \quad (\text{C.6})$$

As a result, (C.3) can be rewritten as

$$\begin{aligned} (\underline{\tilde{\mathbf{J}}}^a \phi)_w \cdot \vec{\nabla}_\xi \phi_h &= (\underline{\tilde{\mathbf{J}}}_1^a)_{rst} \frac{\partial \ell_i^\xi}{\partial \xi} \ell_r^\xi \underbrace{\ell_s^\eta \ell_j^\eta}_{\delta_{sj}} \underbrace{\ell_t^\zeta \ell_k^\zeta}_{\delta_{tk}} \\ &+ (\underline{\tilde{\mathbf{J}}}_2^a)_{rst} \frac{\partial \ell_j^\eta}{\partial \eta} \ell_s^\eta \underbrace{\ell_r^\xi \ell_i^\xi}_{\delta_{ri}} \underbrace{\ell_t^\zeta \ell_k^\zeta}_{\delta_{tk}} \\ &+ (\underline{\tilde{\mathbf{J}}}_3^a)_{rst} \frac{\partial \ell_k^\zeta}{\partial \zeta} \ell_t^\zeta \underbrace{\ell_r^\xi \ell_i^\xi}_{\delta_{ri}} \underbrace{\ell_s^\eta \ell_j^\eta}_{\delta_{sj}}, \end{aligned} \quad (\text{C.7})$$

where δ is Dirac's delta function. Equation (C.7) only takes non-zero values if

$$(s = j \text{ and } t = k) \text{ or } (t = k \text{ and } s = j) \text{ or } (s = j \text{ and } r = i). \quad (\text{C.8})$$

In consequence, there are connectivities between each degree of freedom $h \leftarrow (i, j, k)$ and all degrees of freedom $w \leftarrow (r, s, t)$ that lie along lines of the reference element coordinates. These connectivities appear as non-zero values in the Jacobian matrix, which leads to the following number of non-zeros for the diagonal blocks:

$$nnz_d^a \Big|_{2D} = N_{eq}^2 (P+1)^2 [2(P+1) - 1]. \quad (C.9)$$

$$nnz_d^a \Big|_{3D} = N_{eq}^2 (P+1)^3 [3(P+1) - 2]. \quad (C.10)$$

An entry in the off-diagonal block that connects the degrees of freedom h and w reads [44, 1]

$$\text{ODT}_{hw}^a = \oint_{\partial\Omega \setminus \Gamma} \hat{\mathbf{f}}_{\mathbf{q}^-}^a \phi_w^- \phi_h \, dS, \quad (C.11)$$

706 where ϕ_w^- is the basis function that corresponds to the degree of freedom
707 w , which belongs to an element that is a neighbor of Ω across the interface
708 $\partial\Omega \setminus \Gamma$.

It is evident that ODT_{hw}^a only takes non-zero values if h and w are both degrees of freedom on the boundary $\partial\Omega \setminus \Gamma$. As a result, the number of

non-zero entries for each off-diagonal block reads

$$nnz_o^a \Big|_{2D} = N_{eq}(P+1) \quad (\text{C.12})$$

$$nnz_o^a \Big|_{3D} = N_{eq}(P+1)^2 \quad (\text{C.13})$$

709 *Appendix C.2. Diffusion terms*

Neglecting the advective and time-dependent terms in (C.1), an entry in the diagonal block that connects degrees freedom h and w of a specific element reads [44, 1]

$$\begin{aligned} \text{DT}_{hw}^\nu &= \int_{\Omega} (\underline{\mathbf{J}}^\nu \phi)_w \cdot \vec{\nabla} \phi_h \, d\Omega \\ &+ \sum_{m=1}^{(P+1)^3} \left[\frac{1}{J_m \omega_m} \left(\int_{\Omega} \underline{\mathbf{G}}_m \phi_m \cdot \vec{\nabla} \phi_h \, d\Omega \right) \cdot \left(- \int_{\Omega} \phi_w \vec{\nabla} \phi_m \, d\Omega \right. \right. \\ &\quad \left. \left. + \oint_{\partial\Omega} \hat{\mathbf{q}}_{\mathbf{q}^+} \phi_w \phi_m \vec{n} \, dS + \oint_{\partial\Omega \cap \Gamma} \hat{\mathbf{q}}_{\mathbf{q}^-} \mathbf{q}_{\mathbf{q}^+} \phi_w \phi_m \vec{n} \, dS \right) \right] \\ &\quad - \oint_{\partial\Omega \setminus \Gamma} \left(\hat{\mathbf{f}}_{\mathbf{q}^+}^\nu \phi_w + \hat{\mathbf{f}}_{\vec{\nabla} \mathbf{q}^+}^\nu \cdot \vec{\nabla} \phi_w \right) \phi_h \, dS \\ &\quad - \oint_{\partial\Omega \cap \Gamma} \left(\frac{\partial \hat{\mathbf{f}}_{\Gamma}^\nu}{\partial \mathbf{q}^+} \phi_w + \frac{\partial \hat{\mathbf{f}}_{\Gamma}^\nu}{\partial \vec{\nabla} \mathbf{q}^+} \cdot \vec{\nabla} \phi_w \right) \phi_h \, dS, \quad (\text{C.14}) \end{aligned}$$

where $\underline{\mathbf{J}}^\nu = \partial \hat{\mathbf{f}}^\nu / \partial \mathbf{q}$ is the Jacobian of the diffusive flux with respect to \mathbf{q} , J_m is the Jacobian of the mapping (C.4) at the node m , ω_m are the quadrature weights for the volume integral, $\underline{\mathbf{G}} = \partial \hat{\mathbf{f}}^\nu / \partial (\vec{\nabla} \mathbf{q})$ is the Jacobian of the diffusive flux with respect to $\vec{\nabla} \mathbf{q}$, $\hat{\mathbf{q}}$ is the numerical trace of the solution on the element boundary, $\hat{\mathbf{q}}_{\mathbf{q}^\pm}$ is the derivative of this numerical

trace with respect to the solutions on the element boundary, \vec{n} is the outward-pointing normal vector on the boundary, $\hat{\mathbf{f}}_{\mathbf{q}^+}^\nu$ and $\hat{\mathbf{f}}_{\vec{\nabla}\mathbf{q}^+}^\nu$ are the Jacobians of the viscous surface numerical flux with respect to the solution and its gradient, respectively, and $\partial\hat{\mathbf{f}}_\Gamma^\nu/\partial\mathbf{q}^+$ and $\partial\hat{\mathbf{f}}_\Gamma^\nu/\partial(\vec{\nabla}\mathbf{q}^+)$ are the Jacobians of the viscous surface numerical flux on the physical boundaries. Note that the terms with the subscript Γ contain all the information of the boundary condition on the viscous surface numerical flux:

$$\frac{\partial\hat{\mathbf{f}}_\Gamma^\nu}{\partial\mathbf{q}^+} = \hat{\mathbf{f}}_{\mathbf{q}^+}^\nu + \hat{\mathbf{f}}_{\mathbf{q}^-}^\nu \mathbf{q}_{\mathbf{q}^+}^- + \hat{\mathbf{f}}_{\vec{\nabla}\mathbf{q}^-}^\nu (\vec{\nabla}\mathbf{q}^-)_{\mathbf{q}^+}, \quad \text{and} \quad (\text{C.15})$$

$$\frac{\partial\hat{\mathbf{f}}_\Gamma^\nu}{\partial\vec{\nabla}\mathbf{q}^+} = \hat{\mathbf{f}}_{\vec{\nabla}\mathbf{q}^+}^\nu + \hat{\mathbf{f}}_{\vec{\nabla}\mathbf{q}^-}^\nu (\vec{\nabla}\mathbf{q}^-)_{\vec{\nabla}\mathbf{q}^+}, \quad (\text{C.16})$$

The term first term of the summation in (C.14) is the one that generates the densest sparsity, as it is the multiplication of two volume integrals. This term can be expanded as

$$\begin{aligned} \left(\begin{array}{c} \text{densest} \\ \text{term} \end{array} \right) &= \sum_{m=1}^{(P+1)^3} \left[\frac{1}{J_m \omega_m} \left(\int_{\Omega} \underline{\mathbf{G}}_m \phi_m \cdot \vec{\nabla} \phi_h d\Omega \right) \cdot \left(- \int_{\Omega} \phi_w \vec{\nabla} \phi_m d\Omega \right) \right] \\ &= - \sum_{m=1}^{(P+1)^3} \frac{1}{J_m \omega_m} \left[\left(\int_{\Omega} (\underline{\mathbf{G}}_1 \phi)_m \cdot \vec{\nabla} \phi_h d\Omega \right) \left(\int_{\Omega} \phi_w \frac{\partial \phi_m}{\partial x} d\Omega \right) \right. \\ &\quad + \left(\int_{\Omega} (\underline{\mathbf{G}}_2 \phi)_m \cdot \vec{\nabla} \phi_h d\Omega \right) \left(\int_{\Omega} \phi_w \frac{\partial \phi_m}{\partial y} d\Omega \right) \\ &\quad \left. + \left(\int_{\Omega} (\underline{\mathbf{G}}_3 \phi)_m \cdot \vec{\nabla} \phi_h d\Omega \right) \left(\int_{\Omega} \phi_w \frac{\partial \phi_m}{\partial z} d\Omega \right) \right]. \end{aligned} \quad (\text{C.17})$$

The volume integrals on the left, that depend on the third-order tensors $\underline{\mathbf{G}}_m$, imply two-point connectivities (as in (C.8)) for the degrees of freedom m and h . The volume integrals on the right imply two-point connectivities for the degrees of freedom w and m . In consequence, each degree of freedom $h \leftarrow (i, j, k)$ is connected with non-zeros with all degrees of freedom $w \leftarrow (r, s, t)$ that lie on the same $\xi - \eta$, $\eta - \zeta$ and $\xi - \zeta$ planes of reference element coordinates. Hence, the number of non-zero entries for in the Jacobian matrix in each diagonal block is

$$nnz_d^{\nu} \Big|_{2D} = N_{eq}^2 (P + 1)^4 \quad (\text{C.18})$$

$$nnz_d^{\nu} \Big|_{3D} = 3N_{eq}^2 P (P + 1)^4. \quad (\text{C.19})$$

710 It is important to point out that the sparsity pattern generated by (C.17)
 711 contains all the non-zero entries needed for the other diffusive terms and
 712 for the advective terms. As can be seen, the diffusive terms generate dense
 713 diagonal blocks in 2D.

An entry in the off-diagonal block that connects the degrees of freedom h and w reads [1]

$$\begin{aligned} \text{ODT}_{hw}^{\nu} = \sum_{m=1}^{(P+1)^3} & \left[\frac{1}{J_m \omega_m} \left(\int_{\Omega} \underline{\mathbf{G}}_m \phi_m \cdot \vec{\nabla} \phi_h d\Omega \right) \cdot \left(\int_{\partial\Omega \setminus \Gamma} \phi_w^- \phi_m \vec{n} dS \right) \right] \\ & - \int_{\partial\Omega \setminus \Gamma} \left(\hat{\mathbf{f}}_{\mathbf{q}^-}^{\nu} \phi_w^- + \hat{\mathbf{f}}_{\vec{\nabla} \mathbf{q}^-}^{\nu} \vec{\nabla} \phi_w^- \right) \phi_h dS. \quad (\text{C.20}) \end{aligned}$$

714 In this case, both the summation term and the single surface integral of
 715 (C.20) play an important role in the sparsity of the off-diagonal blocks.

Let us analyse the summation term first. The volume integral implies two-point connectivities for the degrees of freedom m and h , and the surface integral only takes non-zero values if the degrees of freedom w and m lie on an element boundary. As a result, each degree of freedom $h \leftarrow (i, j, k)$ is connected with non-zeros with the degree of freedom (of a neighbor element) $w \leftarrow (r, s, t)$ that lies on the element boundary *and* on the same iso- ξ_i line as h . Therefore, the number of non-zeros due to the summation term is

$$nnz_o^\nu \Big|_{1,2D} = N_{eq}^2 (P+1)^2. \quad (C.21)$$

$$nnz_o^\nu \Big|_{1,3D} = N_{eq}^2 (P+1)^3. \quad (C.22)$$

The single surface integral in (C.20) is important for the sparsity pattern since it contains the gradient of the basis functions on the neighbor element, $\vec{\nabla} \phi_w^-$. This term can be written explicitly as

$$\vec{\nabla} \phi_w^- = \begin{pmatrix} \frac{\partial \phi_w^-}{\partial x} \\ \frac{\partial \phi_w^-}{\partial y} \\ \frac{\partial \phi_w^-}{\partial z} \end{pmatrix} = \begin{pmatrix} \sum_{p=1}^d \frac{\partial \phi_w^-}{\partial \xi_p} \frac{\partial \xi_p}{\partial x} \\ \sum_{p=1}^d \frac{\partial \phi_w^-}{\partial \xi_p} \frac{\partial \xi_p}{\partial y} \\ \sum_{p=1}^d \frac{\partial \phi_w^-}{\partial \xi_p} \frac{\partial \xi_p}{\partial z} \end{pmatrix}. \quad (C.23)$$

Note that the sparsity pattern that this term generates depends on the geometry mapping ($\partial \vec{\xi} / \partial \vec{x}$) and on the position of the degrees of freedom w and h . For a general curvilinear mapping, the second term of (C.20) is zero

when h is not a degree of freedom on the element boundary or when

$$\frac{\partial \phi_w^-}{\partial \xi} = \frac{\partial \phi_w^-}{\partial \eta} = \frac{\partial \phi_w^-}{\partial \zeta} = 0. \quad (\text{C.24})$$

Therefore, for each h on the element boundary, there are non-zeros for the degrees of freedom w of a neighbor element that are arranged along lines of the reference coordinates. In summary, the number of non-zero entries for each off-diagonal block due to the second term of (C.20) is

$$nnz_o^\nu \Big|_{2,2D} = N_{eq}^2 (P+1)[2(P+1) - 1]. \quad (\text{C.25})$$

$$nnz_o^\nu \Big|_{2,3D} = N_{eq}^2 (P+1)^2[3(P+1) - 2]. \quad (\text{C.26})$$

Remark that the term that leads to the non-zero pattern (C.21) shares some non-zeros with the term that leads to (C.25). Combining (C.21) and (C.25), and accounting for the repeated non-zero entries, the total number of non-zeros in an off-diagonal block is

$$\begin{aligned} nnz_o^\nu \Big|_{2D} &= N_{eq}^2 [P(P+1) + (P+1)[2(P+1) - 1]] \\ &= N_{eq}^2 (P+1)(3P+1). \end{aligned} \quad (\text{C.27})$$

$$\begin{aligned} nnz_o^\nu \Big|_{3D} &= N_{eq}^2 [P(P+1)^2 + (P+1)^2[3(P+1) - 2]] \\ &= N_{eq}^2 (P+1)^2(4P+1). \end{aligned} \quad (\text{C.28})$$

716 *Appendix C.3. Total number of non-zero entries*

717 The number of non-zero entries in the diagonal and off-diagonal blocks
 718 depends on the position of the element, i.e. both blocks are more dense
 719 for interior elements connected purely to other interior elements. In our
 720 calculations, we disregard the boundary elements and estimate the upper
 721 bound for the total number of non-zero entries in the Jacobian Matrix:

$$nnz_{full} = N_{el}nnz_d + (C_{Neigh}N_{el} - N_{Out})nnz_o, \quad (C.29)$$

722 where C_{Neigh} is an upper bound of neighbouring elements ($C_{Neigh2D} = 4$
 723 and $C_{Neigh3D} = 6$) and N_{Out} is total number of element faces (3D) or edges
 724 (2D) on the boundary of computational domain. For the cubic mesh used
 725 for Manufactured Solution problem $N_{OutMS3D} = 6(N_{el}^{\frac{1}{3}})^2 = 96$ and for the
 726 NACA0012 case $N_{OutNACA0012} = 880$. The accuracy of these estimations
 727 can be found in Figure C.10. The theoretical curve overestimate the number
 728 of non-zero entries due to the assumptions that were undertaken to estimate
 729 non-zeros in each block and the fact that all the estimated blocks disregard
 730 physical boundary conditions (boundary blocks have significantly less non-
 731 zero entries). The slopes however, follow the same trend within considered
 732 range of polynomials. The reason for the undershoot is twofold. First, the
 733 Jacobian matrices for the Navier-Stokes equations ($\underline{\underline{\mathbf{G}}}$, $\underline{\underline{\mathbf{J}}}^a$ and $\underline{\underline{\mathbf{J}}}^\nu$) are far
 734 from dense (see [44, 1]). Second, the mesh for this case is Cartesian and

735 therefore $\partial\xi_i/\partial x_j = 0$ for $i \neq j$.

736

Now we estimate the number of non-zeros in the condensed system. Due to the two matrix-matrix products (see Section 3.1) needed to compute the Schur complement, the number of non-zero entries in the condensed system significantly increases. The non-zero entries in each block are constrained by the block size, which has complexity (10) ($nb_{bb} = N_{eq}4P$ in 2D and $nb_{bb} = N_{eq}(6P^2 + 2)$ in 3D). However, the SpGEMM operations introduce new non-zero entries into the matrix $\underline{\mathbf{A}}_{cond}$. Additionally, the stencil of the block structure in the Schur complement is wider (non-compact) than in the Jacobian matrix. Therefore, the upper bound for the non-zero entries in the condensed system is


$$3D: \text{nnz}_{cond} = C_{NeighNeigh} N_{el} N_{eq}^2 (6P^2 + 2)^2, \quad (C.30)$$

$$2D: \text{nnz}_{cond} = C_{NeighNeigh} N_{el} N_{eq}^2 (4P)^2, \quad (C.31)$$


737 where the constants $C_{NeighNeigh3D} = 25$ and $C_{NeighNeigh2D} = 13$ place an
 738 upper bound on the total number of blocks per row in the condensed sys-
 739 tem. Note that these constants have been obtained based on the connectives
 740 of structured meshes and can be slightly bigger for particular unstructured
 741 meshes.

742 Finally, Figure C.10 compares the theoretical estimated number of non-
 743 zero entries nnz , for the full and the condensed systems, to the values ex-


744 tracted from the simulations, using the 3D manufactured solution, see section
745 4.1, and the 2D NACA0012 airfoil, see section 4.2. The slopes agree well and
746 it can be seen that the estimates over-predict the simulations in all cases,
747 which follows for having derived upper bounds. Small slope discrepancies for
748 the 2D cases can be explained as follows. Our 2D simulations are not truly
749 2D, but instead we have performed a 3D simulation with only one element in
750 the third direction (and polynomial $P_z = 2$). An approximated upper bound
751 for the nnz (and associated cost) for this particular situation has been ob-
752 tained by assuming three two-dimensional simulations. For this reason when
753 depicting the estimated value in Figure C.10, the estimate has been multi-
754 plied by a constant factor of three. This estimate does not properly account
755 for boundary conditions, which explains small differences.




(a) Manufactured Solution Full System



(b) Manufactured Solution Condensed System



(c) NACA0012 Full System



(d) NACA0012 Condensed System

Figure C.10: Comparison of computed non-zero entries against estimations derived in (C.29), (C.31) and (C.30) for full and condensed systems of two cases considered in this work (Manufactured Solution and NACA0012).

756 **References**

- 757 [1] A. Rueda-Ramírez, E. Ferrer, D. Kopriva, G. Rubio, E. Valero, A
758 statically condensed discontinuous Galerkin spectral element method
759 on Gauss-Lobatto nodes for the compressible Navier-Stokes equations,
760 2019. [arXiv:1911.02366](https://arxiv.org/abs/1911.02366).
- 761 [2] B. Cockburn, C.-W. Shu, The local discontinuous Galerkin method for
762 time-dependent convection-diffusion systems, *SIAM Journal on Numer-*
763 *ical Analysis* 35 (1998) 2440–2463.
- 764 [3] E. Ferrer and R.H.J. Willden, A high order discontinuous Galerkin finite
765 element solver for the incompressible Navier–Stokes equations, *Comput-*
766 *ers & Fluids* 46 (2011) 224–230.
- 767 [4] E. Ferrer, R. H. Willden, A high order discontinuous Galerkin - Fourier
768 incompressible 3D Navier-Stokes solver with rotating sliding meshes,
769 *Journal of Computational Physics* 231 (2012) 7037–7056.
- 770 [5] E. Ferrer, An interior penalty stabilised incompressible discontinuous
771 Galerkin–Fourier solver for implicit large eddy simulations, *Journal of*
772 *Computational Physics* 348 (2017) 754–775.
- 773 [6] N. Fehn, M. Kronbichler, C. Lehrenfeld, G. Lube, P. W. Schroeder,
774 High-order DG solvers for underresolved turbulent incompressible flows:
775 A comparison of L2 and H(div) methods, *International Journal for*
776 *Numerical Methods in Fluids* 91 (2019) 533–556.

- 777 [7] M. Kompenhans, G. Rubio, E. Ferrer, E. Valero, Adaptation strategies
778 for high order discontinuous Galerkin methods based on Tau-estimation,
779 Journal of Computational Physics 306 (2016) 216–236.
- 780 [8] M. Kompenhans, G. Rubio, E. Ferrer, E. Valero, Comparisons of p-
781 adaptation strategies based on truncation- and discretisation-errors for
782 high order discontinuous Galerkin methods, Computers & Fluids 139
783 (2016) 36 – 46. 13th {USNCCM} International Symposium of High-
784 Order Methods for Computational Fluid Dynamics - A special issue
785 dedicated to the 60th birthday of Professor David Kopriva.
- 786 [9] J. Manzanero, E. Ferrer, G. Rubio, E. Valero, Design of a Smagorin-
787 sky spectral Vanishing Viscosity turbulence model for discontinuous
788 Galerkin methods, Computers & Fluids (2020) 104440.
- 789 [10] A. M. Rueda-Ramírez, J. Manzanero, E. Ferrer, G. Rubio, E. Valero,
790 A p-multigrid strategy with anisotropic p-adaptation based on trunca-
791 tion errors for high-order discontinuous Galerkin methods, Journal of
792 Computational Physics 378 (2019) 209–233.
- 793 [11] Z. J. Wang, K. Fidkowski, R. Abgrall, F. Bassi, D. Caraeni, A. Cary,
794 H. Deconinck, R. Hartmann, K. Hillewaert, H. T. Huynh, et al., High-
795 order CFD methods: current status and perspective, International Jour-
796 nal for Numerical Methods in Fluids 72 (2013) 811–845.

- 797 [12] K. Black, A conservative spectral element method for the approximation
798 of compressible fluid flow, *Kybernetika* 35 (1999) 133–146.
- 799 [13] J. Manzanero, G. Rubio, E. Ferrer, E. Valero, D. A. Kopriva, Insights
800 on aliasing driven instabilities for advection equations with application
801 to Gauss–Lobatto discontinuous Galerkin methods, *Journal of Scientific*
802 *Computing* 75 (2018) 1262–1281.
- 803 [14] G. J. Gassner, A. R. Winters, D. A. Kopriva, Split form nodal dis-
804 continuous Galerkin schemes with summation-by-parts property for the
805 compressible Euler equations, *Journal of Computational Physics* 327
806 (2016) 39–66.
- 807 [15] A.R. Winters and G.J. Gassner, Affordable, entropy conserving and
808 entropy stable flux functions for the ideal MHD equations, *Journal of*
809 *Computational Physics* 304 (2016) 72 – 108.
- 810 [16] J. Manzanero, G. Rubio, D. A. Kopriva, E. Ferrer, E. Valero, A
811 free–energy stable nodal discontinuous Galerkin approximation with
812 summation–by–parts property for the Cahn–Hilliard equation, *Journal*
813 *of Computational Physics* 403 (2020) 109072.
- 814 [17] J. Manzanero, G. Rubio, D. A. Kopriva, E. Ferrer, E. Valero, Entropy–
815 stable discontinuous Galerkin approximation with summation–by–parts
816 property for the incompressible Navier–Stokes/Cahn–Hilliard system,
817 *Journal of Computational Physics* (2020) 109363.

- 818 [18] G. J. Gassner, A. R. Winters, F. J. Hindenlang, D. A. Kopriva, The BR1
819 scheme is stable for the compressible Navier–Stokes equations, *Journal*
820 *of Scientific Computing* 77 (2018) 154–200.
- 821 [19] J. Manzanero, G. Rubio, D. A. Kopriva, E. Ferrer, E. Valero, An
822 entropy–stable discontinuous Galerkin approximation for the incom-
823 pressible Navier–Stokes equations with variable density and artificial
824 compressibility, *Journal of Computational Physics* 408 (2020) 109241.
- 825 [20] G. Karniadakis and S.J. Sherwin, *Spectral/hp Element Methods for*
826 *Computational Fluid Dynamics*, Oxford Scholarship, 2005.
- 827 [21] L. Haupt, J. Stiller, W. Nagel, A fast spectral element solver combining
828 static condensation and multigrid techniques, *Journal of Computational*
829 *Physics* 255 (2013) 384 – 395.
- 830 [22] E. Wilson, The static condensation algorithm, *International Journal for*
831 *Numerical Methods in Engineering* 8 (1974) 198–203.
- 832 [23] I. Huismann, J. Stiller, J. Fröhlich, Scaling to the stars – a linearly
833 scaling elliptic solver for p -multigrid, *Journal of Computational Physics*
834 398 (2019) 108868.
- 835 [24] D. Pardo, J. Álvarez Aramberri, M. Paszynski, L. Dalcin, V. Calo, Im-
836 pact of element-level static condensation on iterative solver performance,
837 *Computers and Mathematics with Applications* 70 (2015) 2331–2341.

- 838 [25] S. J. Sherwin, R. M. Kirby, J. Peiró, R. L. Taylor, O. C. Zienkiewicz,
839 On 2D elliptic discontinuous Galerkin methods, *International Journal*
840 *for Numerical Methods in Engineering* 65 (2006) 752–784.
- 841 [26] B. Cockburn, J. Gopalakrishnan, R. Lazarov, Unified hybridization of
842 discontinuous Galerkin, mixed, and continuous Galerkin methods for
843 second order elliptic problems, *SIAM J. Numer. Anal* 47 (2009) 1319–
844 1365. doi:10.1137/070706616.
- 845 [27] J. Carrero, B. Cockburn, D. Schoetzau, Hybridized globally divergence-
846 free LDG methods. part I: The Stokes problem, *Math. Comput.* 75
847 (2006) 533–563. doi:10.1090/S0025-5718-05-01804-1.
- 848 [28] M. Franciolini, K. Fidkowski, A. Crivellini, Efficient discon-
849 tinuous Galerkin implementations and preconditioners for implicit
850 unsteady compressible flow simulations, *arXiv preprint* (2018).
851 [arXiv:physics.comp-ph/1812.04789](https://arxiv.org/abs/physics.comp-ph/1812.04789).
- 852 [29] J. Peraire, N. C. Nguyen, B. Cockburn, An embedded discontinuous
853 Galerkin method for the compressible Euler and Navier-Stokes equa-
854 tions, *20th AIAA Computational Fluid Dynamics Conference 2011*
855 (2011). doi:10.2514/6.2011-3228.
- 856 [30] K. J. Fidkowski, T. A. Oliver, J. Lu, D. L. Darmofal, p-Multigrid so-
857 lution of high-order discontinuous Galerkin discretizations of the com-

- 858 compressible Navier-Stokes equations, *Journal of Computational Physics*
859 207 (2005) 92–113. doi:10.1016/j.jcp.2005.01.005.
- 860 [31] P. O. Persson, An efficient low memory implicit DG algorithm for
861 time dependent problems, *Collection of Technical Papers - 44th AIAA*
862 *Aerospace Sciences Meeting 2* (2006) 1421–1431. doi:10.2514/6.2006-
863 113.
- 864 [32] P. O. Persson, J. Peraire, Newton-GMRES preconditioning for
865 discontinuous Galerkin discretizations of the Navier–Stokes equa-
866 tions, *SIAM Journal on Scientific Computing* 30 (2008) 2709–2733.
867 doi:10.1137/070692108.
- 868 [33] L. T. Diosady, D. L. Darmofal, Preconditioning methods for discontinu-
869 ous Galerkin solutions of the Navier-Stokes equations, *Journal of Com-
870 putational Physics* 228 (2009) 3917–3935. doi:10.1016/j.jcp.2009.02.035.
- 871 [34] K. Shahbazi, D. J. Mavriplis, N. K. Burgess, Multigrid algorithms for
872 high-order discontinuous Galerkin discretizations of the compressible
873 Navier-Stokes equations, *Journal of Computational Physics* 228 (2009)
874 7917–7940. doi:10.1016/j.jcp.2009.07.013.
- 875 [35] P. O. Persson, A sparse and high-order accurate line-based discon-
876 tinuous Galerkin method for unstructured meshes, *Journal of Com-
877 putational Physics* 233 (2013) 414–429. doi:10.1016/j.jcp.2012.09.008.
878 arXiv:1204.1533.

- 879 [36] W. Pazner, P. O. Persson, Stage-parallel fully implicit Runge–Kutta
880 solvers for discontinuous Galerkin fluid simulations, *Journal of Com-*
881 *putational Physics* 335 (2017) 700–717. doi:10.1016/j.jcp.2017.01.050.
882 [arXiv:1701.07181](#).
- 883 [37] M. Franciolini, L. Botti, A. Colombo, A. Crivellini, p-Multigrid matrix-
884 free discontinuous Galerkin solution strategies for the under-resolved
885 simulation of incompressible turbulent flows, 2018. [arXiv:1809.00866](#).
- 886 [38] P. Bastian, E. H. Müller, S. Muthing, M. Piatkowski, Matrix-free multi-
887 grid block-preconditioners for higher order discontinuous Galerkin dis-
888 cretisations, *Journal of Computational Physics* 394 (2019) 417 – 439.
889 doi:<https://doi.org/10.1016/j.jcp.2019.06.001>.
- 890 [39] M. Franciolini, S. M. Murman, Multigrid preconditioning for a space-
891 time spectral-element discontinuous-galerkin solver, *AIAA Scitech 2020*
892 *Forum* (2020). doi:10.2514/6.2020-1314.
- 893 [40] A. Pueyo, D. Zingg, An efficient Newton-GMRES solver for aerodynamic
894 computations, *13th Computational Fluid Dynamics Conference* (1997)
895 712–721. doi:10.2514/6.1997-1955.
- 896 [41] W. Anderson, R. D. Rausch, D. L. Bonhaus, Implicit/multigrid
897 algorithms for incompressible turbulent flows on unstructured
898 grids, *Journal of Computational Physics* 128 (1996) 391 – 408.
899 doi:<https://doi.org/10.1006/jcph.1996.0219>.

- 900 [42] J. Gopalakrishnan, G. Kanschat, A multilevel discontinuous Galerkin
901 method, *Numer. Math.* 95 (2003) 527–550. doi:10.1007/s002110200392.
- 902 [43] K. Black, A conservative spectral element method for the approximation
903 of compressible fluid flow, *Kybernetika* 35 (1999) 133–146.
- 904 [44] A. M. Rueda-Ramírez, Efficient Space and Time Solution Techniques for
905 High-Order Discontinuous Galerkin Discretizations of the 3D Compress-
906 ible Navier-Stokes Equations, Ph.D. thesis, Universidad Politécnica de
907 Madrid, 2019.
- 908 [45] A. Huerta, A. Angeloski, X. Roca, J. Peraire, Efficiency of high-order ele-
909 ments for continuous and discontinuous galerkin methods, *International*
910 *Journal for Numerical Methods in Engineering* 96 (2013) 529–560. URL:
911 <https://onlinelibrary.wiley.com/doi/abs/10.1002/nme.4547>.
912 doi:10.1002/nme.4547. arXiv:<https://onlinelibrary.wiley.com/doi/pdf/10.1002/nme.4547>.
- 913 [46] W. Habchi, *Model Order Reduction (MOR) Techniques*, 2018, pp. 297–
914 338. doi:10.1002/9781119225133.ch8.
- 915 [47] P. Bastian, E. H. Müller, S. Müthing, M. Piatkowski, Matrix-free multi-
916 grid block-preconditioners for higher order discontinuous Galerkin dis-
917 cretisations, *Journal of Computational Physics* 394 (2019) 417 – 439.
- 918 [48] W. Pazner, P. O. Persson, Approximate tensor-product preconditioners
919 for very high order discontinuous Galerkin methods, *Journal of Com-*

- 920 putational Physics 354 (2018) 344–369. doi:10.1016/j.jcp.2017.10.030.
921 arXiv:1704.04549.
- 922 [49] S. Balay, W. D. Gropp, L. C. McInnes, B. F. Smith, Efficient man-
923 agement of parallelism in object oriented numerical software libraries,
924 in: E. Arge, A. M. Bruaset, H. P. Langtangen (Eds.), Modern Software
925 Tools in Scientific Computing, Birkhäuser Press, 1997, pp. 163–202.
- 926 [50] S. Balay, S. Abhyankar, M. F. Adams, J. Brown, P. Brune, K. Buschel-
927 man, L. Dalcin, A. Dener, V. Eijkhout, W. D. Gropp, D. Karpeyev,
928 D. Kaushik, M. G. Knepley, D. A. May, L. C. McInnes, R. T.
929 Mills, T. Munson, K. Rupp, P. Sanan, B. F. Smith, S. Zampini,
930 H. Zhang, H. Zhang, PETSc Users Manual, Technical Report ANL-
931 95/11 - Revision 3.12, Argonne National Laboratory, 2019. URL:
932 <https://www.mcs.anl.gov/petsc>.
- 933 [51] S. Balay, S. Abhyankar, M. F. Adams, J. Brown, P. Brune, K. Buschel-
934 man, L. Dalcin, A. Dener, V. Eijkhout, W. D. Gropp, D. Karpeyev,
935 D. Kaushik, M. G. Knepley, D. A. May, L. C. McInnes, R. T. Mills,
936 T. Munson, K. Rupp, P. Sanan, B. F. Smith, S. Zampini, H. Zhang,
937 H. Zhang, PETSc Web page, <https://www.mcs.anl.gov/petsc>, 2019.
938 URL: <https://www.mcs.anl.gov/petsc>.
- 939 [52] F. Bassi, A. Crivellini, D. A. Di Pietro, S. Rebay, An implicit
940 high-order discontinuous Galerkin method for steady and unsteady

- 941 incompressible flows, *Computers and Fluids* 36 (2007) 1529–1546.
942 doi:10.1016/j.compfluid.2007.03.012.
- 943 [53] B. R. Ahrabi, D. J. Mavriplis, An implicit block ILU smoother for pre-
944 conditioning of Newton–Krylov solvers with application in high-order
945 stabilized finite-element methods, *Computer Methods in Applied Me-
946 chanics and Engineering* 358 (2020) 112637.
- 947 [54] J. Williamson, Low-storage Runge-Kutta schemes, *Journal of Com-
948 putational Physics* 35 (1980) 48 – 56. doi:https://doi.org/10.1016/0021-
949 9991(80)90033-9.
- 950 [55] R. Biswas, K. Devine, J. Flaherty, Parallel, adaptive finite element meth-
951 ods for conservation laws, *Applied Numerical Mathematics* 14 (1994)
952 255–283.
- 953 [56] N. Chalmers, G. Agbaglah, M. Chrust, C. Mavriplis, A parallel hp-
954 adaptive high order discontinuous Galerkin method for the incompress-
955 ible Navier-Stokes equations, *Journal of Computational Physics: X* 2
956 (2019) 100023. doi:https://doi.org/10.1016/j.jcpx.2019.100023.
- 957 [57] P. Birken, G. Gassner, M. Haas, C. D. Munz, Efficient time integra-
958 tion for discontinuous Galerkin method for the unsteady 3D Navier-
959 Stokes equations, *ECCOMAS 2012 - European Congress on Computa-
960 tional Methods in Applied Sciences and Engineering, e-Book Full Papers*
961 (2012) 4334–4353.

- 962 [58] M. J. Zahr, P.-O. Persson, Performance tuning of newton-gmres methods
963 for discontinuous galerkin discretizations of the navier-stokes equations,
964 in: 21st AIAA Computational Fluid Dynamics Conference, 2013, p.
965 2685.
- 966 [59] C. R. Nastase, D. J. Mavriplis, High-order discontinuous Galerkin meth-
967 ods using an hp-multigrid approach, *Journal of Computational Physics*
968 213 (2006) 330–357. doi:10.1016/j.jcp.2005.08.022.
- 969 [60] G. H. Golub, C. F. Van Loan, *Matrix Computations*, third ed., The
970 Johns Hopkins University Press, 1996.
- 971 [61] F. G. Gustavson, Two fast algorithms for sparse matrices: Multiplica-
972 tion and permuted transposition, *ACM Trans. Math. Softw.* 4 (1978)
973 250–269.
- 974 [62] M. Deveci, C. Trott, S. Rajamanickam, Multi-threaded sparse matrix-
975 matrix multiplication for many-core and gpu architectures, *Parallel*
976 *Computing* 78 (2018).
- 977 [63] A. Buluc, J. Gilbert, Parallel sparse matrix-matrix multiplication and
978 indexing: Implementation and experiments, *SIAM Journal on Scientific*
979 *Computing* 34 (2011).
- 980 [64] Y. Saad, M. H. Schultz, GMRES: A Generalized Minimal Residual
981 Algorithm for Solving Nonsymmetric Linear Systems, *SIAM Journal on*
982 *Scientific and Statistical Computing* 7 (1986).

- 983 [65] W. Yang, K. Li, Z. Mo, K. Li, Performance optimization using parti-
984 tioned spmv on gpus and multicore cpus, *IEEE Transactions on Com-*
985 *puters* 64 (2015) 2623–2636.
- 986 [66] T. Vejchodský, P. Šolín, Static condensation, partial orthogonaliza-
987 tion of basis functions, and ILU preconditioning in the hp-FEM, *Journal of Computational and Applied Mathematics* 218 (2008) 192–200.
988 doi:10.1016/j.cam.2007.04.044.
989
- 990 [67] D. A. Kopriva, E. Jimenez, An assessment of the efficiency of nodal
991 discontinuous Galerkin spectral element methods, in: *Recent Develop-*
992 *ments in the Numerics of Nonlinear Hyperbolic Conservation Laws*,
993 Springer, 2013, pp. 223–235.
- 994 [68] I. Huisman, L. Haupt, J. Stiller, J. Fröhlich, Sum factorization of
995 the static condensed Helmholtz equation in a three-dimensional spectral
996 element discretization, *PAMM* 14 (2014). doi:10.1002/pamm.201410465.
- 997 [69] D. A. Kopriva, *Implementing spectral methods for partial differential*
998 *equations: Algorithms for scientists and engineers*, Springer Science &
999 Business Media, 2009.