

1           Impact of Element-Level Static Condensation on  
2                           Iterative Solver Performance

3   D. Pardo<sup>a,b,c</sup>, J. Álvarez-Aramberri<sup>a,b,d</sup>, M. Paszynski<sup>e</sup>, L. Dalcin<sup>f,g</sup>, V. M.  
4                           Calo<sup>g,h</sup>

5   <sup>a</sup>*Department of Applied Mathematics, Statistics, and Operational Research, University of*  
6                           *the Basque Country (UPV/EHU), Bilbao, Spain*

7                           <sup>b</sup>*Basque Center for Applied Mathematics (BCAM), Bilbao, Spain*

8                           <sup>c</sup>*IKERBASQUE, Basque Foundation for Science, Bilbao, Spain*

9                           <sup>d</sup>*University of Pau (UPPA), France*

10                           <sup>e</sup>*AGH University of Science and Technology, Krakow, Poland.*

11                           <sup>f</sup>*National Scientific and Technical Research Council, Argentina*

12                           <sup>g</sup>*Center for Numerical Porous Media, King Abdullah University of Science and*  
13                           *Technology, Saudi Arabia.*

14                           <sup>h</sup>*Applied Mathematics & Computational Science and Earth Science & Engineering, King*  
15                           *Abdullah University of Science and Technology, Saudi Arabia.*

---

16   **Abstract**

This paper provides theoretical estimates that quantify and clarify the savings associated to the use of element-level static condensation as a first step of an iterative solver. These estimates are verified numerically. The numerical evidence shows that static condensation at the element level is beneficial for higher-order methods. For lower-order methods or when the number of iterations required for convergence is low, the setup cost of the elimination as well as its implementation may offset the benefits obtained during the iteration process. However, as the iteration count (e.g., above 50) or the polynomial order (e.g., above cubics) grow, the benefits of element-level static condensation are significant.

17   *Keywords:* Static Condensation, Finite Element Method,  $p$ -FEM, Iterative  
18   Solvers.

---

19   **1. Introduction**

20       Within the Finite Element (FE) community, the term static condensa-  
21       tion of interior degrees of freedom refers to the Gaussian elimination of the

22 element interior *bubble functions* arising from high-order discretizations [Wil-  
23 son (1974)]. Other terms such as Guyan condensation (reduction) can also  
24 be found in the literature to refer to the same set of linear algebra opera-  
25 tions [Guyan (1965)]. Static condensation can also be interpreted as a par-  
26 tial LU factorization of the interior degrees of freedom, as a first step of a  
27 specific substructuring technique, or as a partial orthogonalization of basis  
28 functions [Vejchodsky and Solin (2008)].

29 Interpreted in any of these forms, static condensation constitutes a funda-  
30 mental building block for direct solvers and delivers significant performance  
31 improvements [Calo et al. (2011); Collier et al. (2012, 2014)]. In high-order  
32 methods such as the  $p$ - and  $hp$ -FE methods, interior degrees of freedom are  
33 eliminated first, leading to a reduced system (called Schur complement) that  
34 is subsequently LU factorized. This static condensation step ensures the elim-  
35 ination of interior degrees of freedom before starting the LU factorisation of  
36 the skeleton problem, thereby providing often better performance than that  
37 achieved with traditional ordering techniques, as shown in [Bientinesi et al.  
38 (2010a)]. It also explains why those matrices lacking a structure that enables  
39 static condensation (e.g., higher-continuous basis functions as those used in  
40 meshless methods [Belytschko et al. (1996)], reconstructing kernel element  
41 methods [Collier and Simkins Jr (2009)], and isogeometric analysis [Hughes  
42 et al. (2005)]) typically require a significantly larger number of floating point  
43 operations in order to be factorized [Collier et al. (2012)].

44 While the use of static condensation in direct solvers is always benefi-  
45 cial [Bientinesi et al. (2010b)], its use with iterative solvers is more contro-  
46 versial. Some authors postulate that static condensation should always be  
47 the starting point of any iterative solver, while others refrain from doing so,  
48 since it adds some complexity to the implementation. Even when static con-  
49 densation is unused, most iterative solvers for  $p$ - and  $hp$ -FE methods still  
50 perform some type of elimination (or a spectrally equivalent operation) of lo-  
51 cal interior bubble functions (c.f., [Ainsworth (1996); Pardo and Demkowicz  
52 (2006); Pardo et al. (2006); Beuchler (2013); Arnold et al. (2000); Schoberl  
53 et al. (2008)]).

54 The key point, however, is to determine how profitable it is to explic-  
55 itly build the Schur complement of element bubble functions (as performed  
56 in static condensation) and eliminate the corresponding unknowns from the  
57 global system before performing iterations with respect to keeping the local  
58 LU-factorized matrices as part of the preconditioner without ever computing  
59 the Schur complement. In other words, the distinguishing feature between

60 iterative solvers that employ partial LU factorizations versus those that per-  
61 form static condensation before executing an iterative solver is that the latter  
62 explicitly build the Schur complement and eliminate interior bubble functions  
63 from the global matrix rather than only evaluating their action over a given  
64 residual.

65 This paper provides quantitative estimates about the profitability of us-  
66 ing static condensation before employing an iterative solver. We corroborate  
67 these estimates with numerical experiments in two and three spatial dimen-  
68 sions. Numerical experimentation also enlightens the behavior on the pre-  
69 asymptotic regime. As a result, we describe those situations in which the use  
70 of static condensation is most beneficial. To quantitatively compare both  
71 methods, we present floating point operations (FLOPs) estimates that also  
72 provide interesting clues for the design of optimal hybrid solvers [Soubier  
73 et al. (2011)].

74 In order to make this analysis tractable and easy to follow, we make  
75 several assumptions, which are described in Section 2 along with our model  
76 problem. Section 3 presents precise theoretical complexity estimates illus-  
77 trating the advantages and limitations of using static condensation for each  
78 particular discretization. We describe the implementation details in Section  
79 4 and we present numerical results confirming the estimates in Section 5.  
80 Section 6 describes the conclusions of our study and suggests future research  
81 lines in the topic.

## 82 2. Model Problem and Assumptions

Our starting point is the following algebraic system of linear equations:

$$Ax = b, \tag{1}$$

83 where  $A$  is a non-singular real-valued  $N \times N$  sparse matrix,  $b$  is the right-hand  
84 side, and  $x$  is the solution vector.

85 In this work, we assume that the system matrix  $A$  is associated to a reg-  
86 ular quadrilateral or hexahedral grid coming from a finite element discretiza-  
87 tion with uniform order of approximation  $p$  and with the same number of  
88 elements in each spatial direction. When the number of elements in each  
89 direction is substantially different, then the problem complexity reduces to  
90 that given by a lower dimensional problem.

91 In our estimates and computations, we avoid taking advantage of orthog-  
92 onal basis functions, i.e., we consider all contributions originating from a trial

93 and a test function with shared support as nonzero (a.k.a. “logical nonzero  
 94 entry”), despite the fact that the actual values could indeed be zero. In  
 95 arbitrarily mapped elements (non-affine) and/or in complex bilinear forms,  
 96 logical nonzero entries are indeed different from zero.

97 We assume that the number of iterations needed to solve a given problem  
 98 before static condensation is of the same order as that needed after static  
 99 condensation. A large family of iterative solvers comply with this assumption,  
 100 as shown in the Appendix.

101 We further assume that the cost of building the preconditioner associated  
 102 to the skeleton problem is negligible, since the number of unknowns in the  
 103 skeleton problem is  $\mathcal{O}(p)$  times smaller than those in the interiors of the  
 104 elements. In the case of a multigrid solver, we also assume that the coarse-  
 105 grid correction has a negligible cost, since it consists of solving a smaller-size  
 106 problem.

For simplicity, we restrict our attention to boundary value problems (with  
 Dirichlet boundary conditions) that are governed by second order partial  
 differential equations (PDEs) of the form:

$$\begin{aligned} -\nabla \cdot (c_1 \nabla u) + c_2 \nabla u + c_3 u &= f \quad \text{in } \Omega, \\ u &= u_0 \quad \text{on } \Gamma = \partial\Omega, \end{aligned} \tag{2}$$

107 where  $c_1$  is a symmetric positive definite tensor,  $c_2$  a vector, and  $c_3$  a scalar  
 108 function,  $f$  is the right-hand side,  $u$  is the solution,  $u_0$  is the Dirichlet data,  
 109 and  $\nabla$ ,  $\nabla \cdot$  are the gradient and divergence operators, respectively.  $c_1$ ,  $c_2$ ,  
 110 and  $c_3$  are bounded and spacially varying so they may also incorporate the  
 111 Jacobian of a transformation from the reference elements to a deformed ge-  
 112 ometry [Pardo et al. (2008a,b)]. We also assume that the coefficients are such  
 113 that the above problem has a unique solution.

114 The variational formulation of problem (2) is given by (see e.g., [Hughes  
 115 (2012)]):

$$\begin{cases} \text{Find } u \in u_0 + V \text{ such that,} \\ b(u, v) = l(v) \quad \forall v \in V, \end{cases} \tag{3}$$

where the bilinear form  $b(u, v)$  and the linear form  $l(v)$  are defined as:

$$\begin{aligned}
 b(u, v) &= \int_{\Omega} c_1 \nabla u \cdot \nabla v \, d\Omega + \int_{\Omega} c_2 \nabla u \, v \, d\Omega + \int_{\Omega} c_3 u \, v \, d\Omega \\
 l(v) &= \int_{\Omega} f v \, d\Omega,
 \end{aligned}
 \tag{4}$$

116 with  $v$  a test function belonging to space  $V = H_0^1(\Omega) = \{u \in L^2(\Omega) : u|_{\Gamma} =$   
 117  $0, \nabla u \in \mathbf{L}^2(\Omega)\}$ .

118 For the case of multiple equations and/or  $H(\mathbf{curl})$ ,  $H(\text{div})$ , or  $L^2$  dis-  
 119 cretizations, most of the results presented here can be easily generalized  
 120 using the same methodology as that shown in this paper. For complexity es-  
 121 timates associated to isogeometric analysis (IGA) and finite difference (FD)  
 122 discretizations, we refer to [Collier et al. (2012, 2013)]. In these cases, it  
 123 is not possible to perform static condensation of interior unknowns. Other  
 124 discretizations such as Discontinuous Galerkin (DG) or Hybridizable Discon-  
 125 tinuous Galerkin (HDG) [Nguyen et al. (2011); Kirby et al. (2012)] can be  
 126 analyzed using the same techniques employed here but taking into account  
 127 the precise number of element-interior and element-boundary unknowns de-  
 128 livered by each formulation. For the case of time-domain problems with a  
 129 single time independent matrix, they can be interpreted as a single prob-  
 130 lem with multiple right hand sides, and thus, results shown here can also be  
 131 trivially extended to that situation.

132 Our study focuses on moderate values of  $p$ , namely,  $2 < p < 15$  in 2D  
 133 and  $2 < p < 10$  in 3D. Nonetheless, some tables and graphics showing higher  
 134 values of  $p$  have also been computed for illustration and verification pur-  
 135 poses. For larger values of  $p$ , the costs of integration, memory storage, and  
 136 solution of the system of linear equations may become prohibitively expen-  
 137 sive, and one needs to employ specific spectral method techniques such as  
 138 sum factorization [Vos et al. (2010); Orszag (1980)] or the Spectral Galerkin  
 139 method [Melenk et al. (2001)]. The operation count of both integration and  
 140 matrix-vector multiplication reduces significantly when using either of the  
 141 above techniques for high  $p$ . However, in this paper we consider moderate  
 142 values of  $p$  and we avoid the use of sum factorization techniques. The analysis  
 143 performed in this paper does not consider the use of sparse representations  
 144 of nodes for building the preconditioners [Austin et al. (2012)], which would  
 145 require a separate study.

146 This work focuses on operation counts and does not study paralleliza-

147 tion costs such as communication. However, under the assumption that each  
 148 element is contained only in a single processor (which is typical for the mod-  
 149 erate values of  $p$  considered here), the element-level static condensation is  
 150 performed in a single processor, and no communication costs occur during  
 151 such operation.

### 152 3. Theoretical Complexity Estimates

153 In this section, we derive theoretical complexity estimates to compare  
 154 the number of FLOPs needed to solve a system of linear equations with and  
 155 without static condensation.

156 The considered iterative algorithm consists of the following steps:

- 157 1. For the static-condensation case, construct the Schur complement ma-  
 158 trix in each element.
- 159 2. Construct a block Jacobi preconditioner on the fine grid with blocks as-  
 160 sociated to the unknowns of a single element, as in [Pardo and Demkow-  
 161 icz (2006)]. For the case with static condensation, bubbles have already  
 162 been eliminated in the previous step, thus, blocks are associated to the  
 163 unknowns of the element skeleton.
- 164 3. Construct restriction/prolongation operators and coarse-grid solver for  
 165 the case of multigrid. As described in the previous section, this step  
 166 is assumed to be of lower-order complexity and thus has a negligible  
 167 computational cost.
- 168 4. Iterate (matrix-vector multiplications) until a given error tolerance is  
 169 reached. For the case with static condensation, iterations are performed  
 170 only over the reduced skeleton system.

171 The computational cost associated to the use of other commonly employed  
 172 preconditioners such as Incomplete LU (ILU) factorization with no fill-in —  
 173 ILU(0)— is of the same order, and thus, the analysis performed here also  
 174 applies to those cases.

#### 175 3.1. Notation and Preliminary Considerations

We express the amount of FLOPs needed to solve the above system of  
 equations using an iterative solver without static condensation as:

$$\text{FLOPs} = (n_{it}c + g)M, \tag{5}$$

176 where:

- 177 •  $M$  is the number of nonzero entries in the matrix  $A$ ,
- 178 •  $n_{it}$  is the number of iterations, which depends on the structure of  $A$ ,
- 179 •  $gM$  is the number of FLOPs required to build the preconditioner, and
- 180 •  $c$  is a constant that incorporates the number of matrix-vector multipli-
- 181 cations that need to be performed within each iteration.

When using static condensation, the total amount of FLOPs needed to solve the above system of equations is given by:

$$\text{FLOPs}_{SC} = (\tilde{n}_{it}\tilde{c} + \tilde{g})\tilde{M} + \tilde{s}M, \quad (6)$$

182 where:

- 183 •  $\tilde{M}$  is the number of nonzero entries in the stiffness matrix after static
- 184 condensation,
- 185 •  $\tilde{n}_{it}$  is the number of iterations needed in the statically condensed sys-
- 186 tem,
- 187 •  $\tilde{g}\tilde{M}$  is the number of FLOPs required to build the preconditioner of
- 188 the statically condensed system,
- 189 •  $\tilde{c}$  is a constant that incorporates the number of matrix-vector multi-
- 190 plications performed within each iteration of the statically condensed
- 191 system, and
- 192 •  $\tilde{s}M$  is the number of FLOPs needed to perform static condensation
- 193 over the original system (independent of  $\tilde{n}_{it}$ ).

194 The above formulas depend upon the order of approximation  $p$ . In other  
 195 words,  $g$ ,  $M$ ,  $\tilde{g}$ ,  $\tilde{M}$ , and  $s$  are  $p$  dependent. Additionally,  $n_{it}$ , and  $\tilde{n}_{it}$  may  
 196 depend upon  $p$ , but in here we assume that the preconditioner is able to keep  
 197 the number of iterations almost constant irrespective of  $p$ . In fact, there exist  
 198 several iterative solvers in the literature for which the number of iterations  
 199 needed to converge is almost independent of  $p$  (c.f., [Ainsworth (1996); Pardo  
 200 and Demkowicz (2006); Schoberl et al. (2008); Beuchler (2013); Gao (2013)]).

201 The objective of static condensation is to reduce the total cost by reducing  
 202 the value of the constant  $\tilde{M}$  versus  $M$ , possibly at the cost of some small

203 overhead  $\tilde{s}M$  needed to build the Schur complement matrix. In other words,  
 204 our analysis focuses on the amount of operations needed to build the Schur  
 205 complement and the savings in terms of number of nonzero entries that this  
 206 system may deliver during iterations (matrix-vector multiplications).

207 The speedup factor in terms of the number of FLOPs due to the use of  
 208 static condensation is independent of the number of elements (up to bound-  
 209 ary conditions). In order to properly normalize the results, we provide all  
 210 estimates on a per-element basis. We assume that a sufficiently large number  
 211 of elements is used, thus making the effect of the boundaries irrelevant.

212 To compute the number of independent nodes per element for a problem  
 213 with infinite number of elements in each direction, we consider a single hex-  
 214 ahedron problem with periodic boundary conditions [Calo et al. (2011); Col-  
 215 lier et al. (2012)]. The element has four different types of nodes (`node_type`),  
 216 namely, vertices, edges, faces, and volumetric interiors. The number of nodes  
 217 of each particular type—`nr_nodes(node_type)`—in an element with periodic  
 boundary conditions is displayed in Table 1, and illustrated in Figure 3.1.

Denomination	Dimension ( $d$ )	Nr_Vert	Nr_Edg	Nr_Fac	Nr_Vol
2D	2	1	2	1	0
3D	3	1	3	3	1

Table 1: Number of independent vertices (Nr\_Vert), edges (Nr\_Edg), faces (Nr\_Fac), and volumes (Nr\_Vol) for a single element (quadrilateral in 2D or hexahedron in 3D) with periodic boundary conditions.

218

The problem size (per element)  $N$  is given by:

$$N = \sum_{\text{node\_type}} \text{nr\_nodes}(\text{node\_type}) \cdot \text{dof}(\text{node\_type}), \quad (7)$$

where `dof(node_type)` is the number of degrees of freedom for each `node_type`. Let  $e$  be the node dimension of `node_type`, i.e., 0 for vertices, 1 for edges, 2 for faces, and 3 for volumes. For an  $H^1(C^0)$  element of order  $p$ , the number of degrees of freedom is equal to

$$\text{dof}(\text{node\_type}) = (p - 1)^e. \quad (8)$$






Figure 1: Quadrilateral (left panel) and Hexahedral (right panel) in a periodic mesh. V=Vertex, E=Edge, F=Face. PV, PE, and PV correspond to vertices, edges and faces where periodic boundary conditions are imposed, and they are denoted with the red color.

219 *3.2. Without Static Condensation*

220 We first consider the case of an iterative solver built without performing  
 221 static condensation. Following equation (5), we decompose the total cost into  
 222 two parts: (a) building the preconditioner, and (b) performing iterations.

223 *Cost of Building the Preconditioner.* For the particular case of a multigrid  
 224 solver, we assume that the construction of the transfer operators, smoothers,  
 225 and the coarsest grid solve are significantly cheaper operations.

226 *Cost of Iterations.* The number of FLOPs needed to perform each iteration  
 227 is dominated by two matrix-vector multiplications: one corresponding to the  
 228 original matrix  $A$ , and the second one corresponding to the preconditioner  
 229 matrix. Assuming that the number of nonzero entries in the preconditioner  
 230 is smaller or proportional to that of the original matrix, we conclude that the  
 231 iteration cost is proportional to the number of nonzero entries in the original  
 232 matrix.

Then, for an infinitely large problem, the number of nonzero entries per element in the original system  $M$  is given by:

$$M = \sum_{\text{node\_type}} \text{nr\_nodes}(\text{node\_type}) \cdot \text{dof}(\text{node\_type}) \cdot \text{interact\_dof}(\text{node\_type}), \quad (9)$$

where `interact_dof(node_type)` is the number of degrees of freedom interacting with a given `node_type`, and its formula is given by:

$$\text{interact\_dof}(\text{node\_type}) = (2p + 1)^{d-e} \cdot (p + 1)^e. \quad (10)$$

### 233 3.3. With Static Condensation

234 In this case, we need to first consider the number of FLOPs needed to  
 235 perform static condensation. Additionally, we also need to consider the com-  
 236 plexity of building the preconditioner and to perform the iterations.

*Cost of Static Condensation.* Let matrix  $A$  be decomposed as:

$$A = \begin{bmatrix} A_{II} & A_{IS} \\ A_{SI} & A_{SS} \end{bmatrix}, \quad (11)$$

where subscript  $I$  corresponds to the interior bubble degrees of freedom to be eliminated from the system by static condensation and  $S$  to the remaining skeleton unknowns. Performing the partial LU factorization of the square submatrix  $A_{II}$ , we obtain:

$$A = \begin{bmatrix} \mathbf{I} & 0 \\ A_{SI}A_{II}^{-1} & \mathbf{I} \end{bmatrix} \cdot \begin{bmatrix} A_{II} & 0 \\ 0 & A_{SS} - A_{SI}A_{II}^{-1}A_{IS} \end{bmatrix} \cdot \begin{bmatrix} \mathbf{I} & A_{II}^{-1}A_{IS} \\ 0 & \mathbf{I} \end{bmatrix}, \quad (12)$$

237 where  $S = A_{SS} - A_{SI}A_{II}^{-1}A_{IS}$  is the Schur complement. In practice, a full  
 238 inversion of  $A_{II}$  should be avoided and replaced by the corresponding LU fac-  
 239 torization. First, because round-off errors are less relevant when performing  
 240 an LU factorization. Second, because the number of FLOPs also diminishes,  
 241 specially, for high values of  $p$ . Although both operations (LU factorization  
 242 and inversion) exhibit the same scaling in terms of  $p$ , the corresponding con-  
 243 stants are different. More precisely, inverting the matrix is about three times  
 244 more expensive than performing its LU factorization.<sup>1</sup>

---

<sup>1</sup>Computing the inverse of a matrix  $A$  of size  $n \times n$  with LAPACK [Anderson et al. (1999)] requires two steps: first, compute the factorization  $A = LU$  (by calling `DGETRF`), next compute  $A^{-1}$  out of the LU factors (by calling `DGETRI`). The LU factorization of a matrix of size  $n \times n$  has a flop count of  $\mathcal{O}(\frac{2}{3}n^3)$  and the computation of the inverse out for the LU factors has a flop count of  $\mathcal{O}(\frac{4}{3}n^3)$  (see [Blackford and Dongarra (1999)], pp. 120). Therefore, the computation of the matrix inverse has a total flop count of  $\mathcal{O}(2n^3)$ , which is three times the one required to compute the LU factors.

245 Taking the above observation into account, we conclude that the cost of  
 246 building the Schur complement (per element) is the sum of the costs asso-  
 247 ciated to: (a) performing the  $LU$  factorization of matrix  $A_{II}$ , (b) comput-  
 248 ing the action  $A_{II}^{-1}A_{IS}$  by using backward and forward substitutions, and  
 249 (c) performing a matrix-matrix multiplication. The cost of performing the  
 250  $LU$  factorization of matrix  $A_{II}$  is of order  $\mathcal{O}(p^{3d})$ , while the cost of perform-  
 251 ing backward and forward substitutions is of order  $\mathcal{O}(p^{3d-1})$ . The remaining  
 252 operations are of lower order, thus their contribution to the total number of  
 253 FLOPs is negligible.

254 *Cost of Building the Preconditioner.* As stated in previous sections, we as-  
 255 sume that the cost of building the preconditioner for the skeleton problem  
 256 is negligible, since the number of unknowns of this problem is  $\mathcal{O}(p)$  times  
 257 smaller than that of the original problem.

258 *Cost of Iterations.* As in the previous case, we assume that the number  
 259 of FLOPs is proportional to the number of nonzero entries in the Schur  
 260 complement system  $S$ .

If we denote the nodes of type “skeleton” (all except the volumetric inte-  
 rior ones) by `—node_ske—`, then, for an infinitely large problem the number  
 of nonzero entries per element in the statically condensed system  $\tilde{M}$  is given  
 by:

$$\begin{aligned} \tilde{M} = \sum_{\text{node\_ske}} \text{nr\_nodes}(\text{node\_ske}) \times \text{dof}(\text{node\_ske}) \\ \times [\text{interact\_dof}(\text{node\_ske}) - 2^{d-e}(p-1)^d]. \end{aligned} \quad (13)$$

261 In the above formula, the term “`interact_dof (node_ske) - 2d-e(p - 1)d”` cor-  
 262 responds to all interacting nodes of the given “`node_ske`” minus those corre-  
 263 sponding to volumetric interior nodes already eliminated by the static con-  
 264 densation.

A simple asymptotic expansion shows that  $M = \mathcal{O}(p^{2d})$ , while  $\tilde{M} = \mathcal{O}(p^{2(d-1)})$ , which implies that the number of nonzero entries in the statically condensed system is  $\mathcal{O}(p^2)$  times smaller than in the original system. Indeed, we can further derive the following approximations for the number of nonzero

entries per element, whose accuracy are confirmed by Tables 2 and 3:

$$M(2D) \approx (p+1)^3 p \qquad M(3D) \approx (p+1)^5 p \qquad (14)$$

$$\tilde{M}(2D) \approx 14p^2 \qquad \tilde{M}(3D) \approx 33p^4 \qquad (15)$$

Table 2: 2D results for a problem with periodic boundary conditions (infinitely large problem). Number of nonzero entries per element with and without static condensation.

$p$	1	2	3	4	5	10	25	100	1000
$M/(p+1)^3 p$	1.12	1.19	1.17	1.15	1.13	1.08	1.04	1.01	1.00
$\tilde{M}/14p^2$	0.64	0.84	0.90	0.92	0.94	0.97	0.99	1.00	1.00
$M/\tilde{M}$	1	1.36	1.99	2.78	3.72	10.6	52.7	745	71735

Table 3: 3D results for a problem with periodic boundary conditions (infinitely large problem). Number of nonzero entries per element with and without static condensation.

$p$	1	2	3	4	5	10	25	100	1000
$M/((p+1)^5 p)$	0.84	1.05	1.10	1.11	1.10	1.07	1.04	1.00	1.00
$\tilde{M}/(33p^4)$	0.82	0.87	0.90	0.92	0.94	0.97	0.99	1.00	1.00
$M/\tilde{M}$	1	1.11	1.40	1.77	2.22	5.42	24.2	322	30496

### 265 3.4. Discussion and Final Estimates

Taking into account all the above estimates, the number of FLOPs needed to solve the original and the statically condensed systems are respectively of the order:

$$\begin{aligned} \text{FLOPs} &= \mathcal{O}(n_{it} p^{2d}) + \mathcal{O}(p^{3d}), \\ \text{FLOPs}_{SC} &= \mathcal{O}(n_{it} p^{2(d-1)}) + \mathcal{O}(p^{3d}) + \mathcal{O}(p^{3d-1}), \end{aligned} \qquad (16)$$

266 per element.

267 Although  $\mathcal{O}(p^{3d-1})$  can be considered as a lower order term in the above  
 268 estimates, it sometimes becomes significant for low and moderate values of  $p$ ,  
 269 since the constant in front of it is significantly higher than that of the higher  
 270 order term  $\mathcal{O}(p^{3d})$ .

271 In any case, we observe that as  $p \rightarrow \infty$ , the number of FLOPs is the  
 272 same regardless of the use of static condensation. Thus, for sufficiently high

273  $p$ , perhaps the additional burden of implementing static condensation should  
274 be avoided, and the main effort should be focused on performing some type  
275 of incomplete LU factorization (or some alternative iterative solver) on the  
276 bubble degrees of freedom.

277 On the other hand, for a fixed value of  $p$ , we see that as  $n_{it} \rightarrow \infty$ ,  
278 the use of static condensation is always recommended. To determine the  
279 exact regions for which the use of static condensation brings any benefit, we  
280 complement the above theoretical estimates with numerical experimentation.  
281 These simulation results allow us to estimate the relative significance of the  
282 different contributions of each term in (16).

## 283 4. Implementation

284 To provide further insight to the above theoretical estimates, we built a  
285 FE code in one-, two-, and three-spatial dimensions. The method is imple-  
286 mented in MATLAB, and supports any uniform order of approximation  $p$ .  
287 For its construction, we have assumed a tensor product structure for the basis  
288 functions and the corresponding discretization. To simplify the implementa-  
289 tion, the size of all elements in each direction is selected to be equal to one,  
290 and we have chosen  $c_1 = 1$ ,  $c_2 = 0$ , and  $c_3 = 0$  everywhere, which produces a  
291 constant Jacobian (in fact, equal to one) and facilitates the decomposition of  
292 the bilinear form as the sum of tensor products of one-dimensional problems.  
293 Using this identity map for the geometry does not affect the nonzero pattern  
294 of the system, while keeping the implementation simple.

## 295 5. Numerical Results

296 In this section, we solve the Laplace equation (2) with Dirichlet bound-  
297 ary conditions using the FE method, as described in the previous section.  
298 Numerical results are obtained in a workstation equipped with 94GB RAM,  
299 and an eight-core Intel Xeon E5620 processor (12Mb cache, 2.40 Ghz).

### 300 5.1. Number of Nonzero Entries

301 We first report the number of nonzero entries per element for the con-  
302 densed and non-condensed systems. In Table 4, we observe that the nu-  
303 merical results for  $2D$  match with the estimates of Table 2. For example,  
304 for  $p = 5$ , the number of nonzero entries for the statically condensed and  
305 non-condensed systems computed numerically coincide with the theoretical

306 estimates (1.13 and 0.94, respectively). Their ratio is also the same (3.73 for  
 307 the numerical value and 3.72 for the estimate).

308 In the 3D case (Table 5), the agreement with the estimates (Table 3)  
 309 is more visible for low  $p$ . Estimates correspond to the case of an infinitely  
 310 large problem. Therefore, when we solve the Laplace equation with Dirichlet  
 311 boundary conditions, in order to obtain a perfect match between the nu-  
 312 merical results and the estimates, we need to consider a sufficiently large  
 313 problem. For small polynomial orders we are able to solve large problems.  
 314 This is reflected in the existing good agreement between numerical results  
 315 and theoretical estimates (e.g., for  $p = 3$ ). However, an increase in  $p$  limits us  
 316 from considering as many elements as needed to observe a perfect agreement  
 317 (e.g.,  $p = 10$ ).

318 Tables 6 and 7 illustrate that when the the problem size increases, the  
 319 numerical results converge to the theoretical estimates, as predicted.

Table 4: 2D numerical results for a problem with a number of elements equal to  $1000^2$ ,  $333^2$ ,  $200^2$ ,  $100^2$ ,  $40^2$  and  $20^2$  for  $p$  equal to 1, 3, 5, 10, 25 and 40, respectively. Number of nonzero entries per element with and without static condensation.

$p$	$p = 1$	$p = 3$	$p = 5$	$p = 10$	$p = 25$	$p = 40$
$M/(p+1)^3 p$	1.12	1.17	1.13	1.08	1.03	1.02
$\tilde{M}/14p^2$	0.64	0.89	0.94	0.96	0.97	0.96
$M/\tilde{M}$	1	1.99	3.73	10.64	53.45	130.98

Table 5: 3D numerical results for a problem with a number of elements equal to  $100^3$ ,  $33^3$ ,  $20^3$  and  $8^3$  for  $p$  equal to 1, 3, 5, and 10, respectively. Number of nonzero entries per element with and without static condensation.

$p$	$p = 1$	$p = 3$	$p = 5$	$p = 10$
$M/(p+1)^5 p$	0.84	1.07	1.06	1.01
$\tilde{M}/(33p^4)$	0.81	0.87	0.89	0.86
$M/\tilde{M}$	1	1.40	2.24	5.74

## 320 5.2. Time Comparison

Assuming that the amount of time needed to execute an algorithm is proportional to the number of FLOPs, we can reinterpret equations (5) and (6) in terms of time instead of FLOPs by rewriting them in the following way:

$$TIME = T(A \cdot \bar{v})n_{it} + T(A_{II}), \quad (17)$$

Table 6: 2D numerical results showing how the number of nonzero entries, with and without static condensation, converges to the theoretical estimates when the number of elements ( $N_e$ ) increases.

$p = 5$	$N_e = 5^2$	$N_e = 25^2$	$N_e = 50^2$	$N_e = 100^2$	$N_e = 200^2$	Estimate
$M/(p+1)^3p$	1.02	1.11	1.12	1.13	1.13	1.13
$\tilde{M}/14p^2$	0.81	0.91	0.93	0.93	0.94	0.94
$M/\tilde{M}$	3.91	3.76	3.74	3.73	3.73	3.72

Table 7: 3D numerical results showing how the number of nonzero entries, with and without static condensation, converges to the theoretical estimates when the number of elements ( $N_e$ ) increase.

$p = 5$	$N_e = 3^3$	$N_e = 5^3$	$N_e = 10^3$	$N_e = 20^3$	Estimate
$M/(p+1)^5p$	0.84	0.94	1.02	1.06	1.10
$\tilde{M}/(33p^4)$	0.66	0.76	0.85	0.89	0.94
$M/\tilde{M}$	2.41	2.33	2.27	2.24	2.22

and

$$TIME_{SC} = T(A_{SS} \cdot \bar{v})n_{it} + T(A_{SS}), \quad (18)$$

321 where  $TIME$  and  $TIME_{SC}$  are the total times without and with static  
322 condensation, respectively. Here,  $n_{it}$  is again the number of iterations the  
323 iterative solver requires to converge,  $T(A_{SS})$  is the time required to compute  
324 the Schur complement,  $T(A_{II})$  corresponds to the time devoted to perform  
325 the LU factorization of  $A_{II}$ , and  $T(A \cdot \bar{v})$  and  $T(A_{SS} \cdot \bar{v})$  are, respectively,  
326 the time required to multiply the non-condensed stiffness matrix  $A$  and the  
327 statically condensed stiffness matrix  $A_{SS}$  by a vector. As before, we assume  
328 that the time needed to build the preconditioner in the condensed system is  
329 negligible.

330 Since these times scale linearly with the number of elements, all numerical  
331 results reported in Tables 8 and 9 are given in terms of time (seconds) per  
332 element.

333 According to Section 3.3, for a sufficiently large  $p$ , the dominant part  
334 in terms of number of FLOPs is the LU factorization of  $A_{II}$ , regardless the  
335 decision of using static condensation or not. When the Schur complement is  
336 explicitly computed, additional backward and forward substitutions are re-  
337 quired. Since the ratio between interior and skeleton nodes tends to infinity  
338 as  $p$  grows, the cost of these forward and backward substitutions becomes

339 negligible for sufficiently large  $p$ , and the quotient  $T(A_{SS})/T(A_{II})$  approaches  
 340 1. However, for small values of  $p$ , the situation may vary considerably, since  
 341 the number of interior and skeleton nodes are comparable. This is numerically  
 342 observed in Tables 8 and 9, where we also appreciate that the ratio  
 343  $T(A_{SS})/T(A_{II})$  tends to 1 for large  $p$ , as expected.

344 We also know that the time required to multiply a matrix times a vector  
 345 depends on the number of nonzeros in the matrix. Therefore, we expect to  
 346 observe that the ratios  $T(A \cdot \bar{v})/T(A_{SS} \cdot \bar{v})$  and  $M/\tilde{M}$  behave in a similar  
 347 fashion. This is effectively observed for example by comparing the last row of  
 348 Table 9 (with values equal to 1.41, 2.32 and 5.77) with the last row of Table 5  
 349 (with values equal to 1.40, 2.24 and 5.74, respectively).

Table 8: 2D numerical results for a problem with number of elements equal to  $33^2$ ,  $20^2$ ,  $10^2$ ,  $4^2$ ,  $2^2$  for  $p$  equal to 3, 5, 10, 25, 40 and 100, respectively. Times per element with and without static condensation. Times have been computed 100 times and averaged.

$p$	$p = 3$	$p = 5$	$p = 10$	$p = 25$	$p = 40$	$p = 100$
$T(A_{SS})$	1.74e-5	3.92e-5	7.02e-4	2.31e-2	2.23e-1	16.94
$T(A_{II})$	2.95e-6	1.08e-5	2.09e-4	1.18e-2	1.55e-1	14.79
$T(A_{SS} \cdot \bar{v})$	3.13e-7	8.63e-7	3.26e-6	2.42e-5	5.11e-5	2.13e-4
$T(A \cdot \bar{v})$	6.21e-7	3.56e-6	3.51e-5	1.05e-3	6.63e-3	2.55e-1
$T(A_{SS})/T(A_{II})$	5.90	3.63	3.36	1.96	1.44	1.15
$T(A \cdot \bar{v})/T(A_{SS} \cdot \bar{v})$	1.98	4.13	10.77	43.39	129.75	1197.2

Table 9: 3D numerical results for a problem with number of elements equal to  $33^3$ ,  $20^3$ ,  $8^3$ ,  $2^3$  and 1 for  $p$  equal to 3, 5, 10, 15 and 25, respectively. Times per element with and without static condensation. Times have been computed 100 times and averaged.

$p$	$p = 3$	$p = 5$	$p = 10$	$p = 15$	$p = 25$
$T(A_{SS})$	4.34e-5	7.48e-4	7.60e-2	1.49	77.76
$T(A_{II})$	4.13e-6	8.76e-5	2.09e-2	5.30e-1	36.85
$T(A_{SS} \cdot \bar{v})$	5.58e-6	4.19e-5	7.11e-4	2.34e-3	1.02e-2
$T(A \cdot \bar{v})$	7.89e-6	9.70e-5	4.10e-3	3.30e-2	6.23e-1
$T(A_{SS})/T(A_{II})$	10.51	8.54	3.64	2.81	2.11
$T(A \cdot \bar{v})/T(A_{SS} \cdot \bar{v})$	1.41	2.32	5.77	14.10	61.08

350 Figure 2 displays the ratio between the time required to solve a 2D prob-



351 lem without static condensation against that needed to solve the same prob-  
 352 lem with static condensation. As predicted by the theory, the use of static  
 353 condensation is always recommended for a sufficiently large number of iter-  
 354 ations. As the number of iterations grows, the cost of matrix-vector multi-  
 355 plications becomes dominant, and the savings are of the order of  $p^2$ , which  
 356 correspond to  $T(A \cdot \bar{v})/T(A_{SS} \cdot \bar{v})$  row in Tables 8 and 9. In the pre-asymptotic  
 357 regime, these savings become substantial when the number of iterations is  
 358 above 100. On the other hand, when the number of iterations is below 10,  
 359 the use of static condensation provides no advantage. We also observe that  
 360 as the value of  $p$  increases, the number of iterations needed for static con-  
 361 densation to be beneficial diminishes. Similar conclusions can be depicted from  
 362 the 3D results displayed in Figure 3.




Figure 2: Ratio between the time needed to solve a 2D problem without static condensation vs. that using static condensation.

## 363 6. Conclusions

364 The use of static condensation in iterative solvers is controversial. That  
 365 is, some authors defend its use as an starting point of any iterative solver,




Figure 3: Ratio between the time needed to solve a 3D problem without static condensation vs. that using static condensation.

366 while others do not. The key point is to determine whether computing the  
 367 Schur complement instead of keeping the local LU-factorized matrices as a  
 368 preconditioner is profitable or not.

369 In the present work, we provide quantitative theoretical estimates to de-  
 370 termine the situations where it is recommended to use the static condensa-  
 371 tion, and illustrated these estimates with numerical experiments. To make  
 372 the problem tractable, we have made several assumptions on the mesh regu-  
 373 larity and on the behavior of the iterative solver so the number of iterations  
 374 becomes independent of the mesh size. Under these assumptions, we show  
 375 that the use of static condensation is most beneficial when the number of iter-  
 376 ations is sufficiently large, assuming that  $p > 3$ . However, when the number  
 377 of iterations is below 10, the use of static condensation may be counterpro-  
 378 ductive.

379 Therefore, a more nuanced view point is necessary. Instead of using al-  
 380 ways this technique as a starting point for an iterative solver, we recommend  
 381 to use the tables provided in this paper to analyze the profitability of this  
 382 technique when dealing with high-order methods. In particular, if the iter-

383 ative solver converges in over 50 iterations, the use of static condensation is  
384 highly recommended.

385 This work also provides a first step towards obtaining rigorous computa-  
386 tional complexity estimates for hybrid solvers. Based on existing estimates  
387 for uniform and highly refined grids, the construction of further Schur com-  
388 plements (corresponding to a group of elements) before employing an itera-  
389 tive solver seems inadequate for uniform grids, however, quite profitable and  
390 advantageous for certain types of highly refined grids. This will be analyzed  
391 in detail in a future contribution.

## 392 **Appendix A. Effect of Static Condensation on the Required Num- 393 ber of Iterations**

394 In general, the use of element-level static condensation reduces the condi-  
395 tion number of the original matrix in terms of polynomial order of approxi-  
396 mation  $p$ , although not in terms of element size  $h$  (see, e.g., [Vejchodsky and  
397 Solin (2008)]). However, this positive effect on the condition number of the  
398 unpreconditioned stiffness matrix often has little relevance, since the actual  
399 objective is to reduce the condition number of the *preconditioned* system.

400 For poorly designed solvers (e.g., those using no preconditioner or only  
401 a diagonal preconditioner), the use of static condensation may significantly  
402 reduce the number of iterations required to achieve a given tolerance error.  
403 In such cases, our results need to be modified to account for the savings on  
404 the number of iterations produced by the use of static condensation.

405 However, for well-designed solvers that incorporate proper precondition-  
406 ers in terms of  $p$ , the savings on the number of iterations produced by the  
407 use of static condensation may vanish. For example, any ILU preconditioner  
408 (when properly ordered by considering first the element interior unknowns)  
409 provides a preconditioned system whose condition number is independent  
410 of the use of static condensation, since one obtains the same algebraic sys-  
411 tem [Vejchodsky and Solin (2008)] for both cases. Indeed, such precondition-  
412 er is optimal for dealing with the element interior unknowns, since it is  
413 equivalent to a full LU factorization on the element interiors, because the  
414 solution of interior unknowns introduces no additional fill in.

415 Another prominent family of preconditioners that require the same num-  
416 ber of iterations to converge irrespectively of the use (or not) of static conden-  
417 sation when combined with CG or GMRES is an overlapping block Jacobi  
418 smoother with blocks determined by those basis functions whose support

419 is fully contained in the support of a given vertex basis function (see, for  
 420 instance, [Schoberl et al. (2007), Pardo (2004), page 84]). Numerical exper-  
 421 iments confirm this result, as illustrated on Table A.10. In this table, the  
 422 number of iterations required to converge without and with static condensa-  
 423 tion are almost identical in all cases. These numerical results also confirm  
 424 that static condensation only improves the condition number with respect to  
 425  $p$ , and not  $h$ . However, an adequate preconditioner in  $p$  will achieve a similar  
 426 effect.

Table A.10: Number of iterations required to reduce the residual by five orders of magni-  
 tude without (first number within each cell) and with (second number within each cell)  
 the use of static condensation for a 2D Laplace problem using an overlapping Block-Jacobi  
 smoother [Pardo (2004)].

Order of approximation	$16 \times 16$ elements	$32 \times 32$ elements	$64 \times 64$ elements
$p = 2$	23/23	43/43	87/85
$p = 4$	22/22	43/43	87/86
$p = 6$	22/22	44/43	87/86
$p = 8$	23/23	44/44	88/86

427 When the above preconditioners are combined with a multigrid solver,  
 428 the convergence properties also remain independent with respect to the use  
 429 or not of static condensation.

### 430 **Acknowledgment**

431 This work was partially supported by the European Union’s Horizon  
 432 2020 research and innovation programme under the Marie Skłodowska-Curie  
 433 grant agreement No 644602, the Project of the Spanish Ministry of Econ-  
 434 omy and Competitiveness with reference MTM2013-40824-P, the BCAM  
 435 “Severo Ochoa” accreditation of excellence SEV-2013-0323, the CYTED 2011  
 436 project 712RT0449, the Basque Government Consolidated Research Group  
 437 Grant IT649-13 on “Mathematical Modeling, Simulation, and Industrial Ap-  
 438 plications (M2SI)”, the Polish National Science Center grant no. DEC-  
 439 2012/06/M/ST1/00363 and the Center for Numerical Porous Media at King  
 440 Abdullah University of Science and Technology (KAUST). Lisandro Dalcin  
 441 was partially supported by Agencia Nacional de Promoción Científica y Tec-  
 442 nológica grants PICT 2014-2660 and PICT-E 2014-0191. This publication

443 also was made possible by a National Priorities Research Program grant from  
444 the Qatar National Research Fund (a member of The Qatar Foundation).  
445 The statements made herein are solely the responsibility of the authors.

## 446 **Bibliography**

447 Ainsworth, M., 1996. A preconditioner based on domain decomposition for  
448 h-p finite-element approximation on quasi-uniform meshes. *SIAM Journal*  
449 *on Numerical Analysis* 33 (4), 1358–1376.

450 Anderson, E., Bai, Z., Bischof, C., Blackford, S., Demmel, J., Dongarra, J.,  
451 Du Croz, J., Greenbaum, A., Hammarling, S., McKenney, A., Sorensen,  
452 D., 1999. *LAPACK Users' Guide*, 3rd Edition. Society for Industrial and  
453 Applied Mathematics, Philadelphia, PA.

454 Arnold, D. N., Falk, R. S., Winther, R., 2000. Multigrid in  $H(\text{div})$  and  
455  $H(\text{curl})$ . *Numer. Math.* 85 (2), 197–217.

456 Austin, T. M., Brezina, M., Jamroz, B., Jhurani, C., Manteuffel, T. A., Ruge,  
457 J., 2012. Semi-automatic sparse preconditioners for high-order finite ele-  
458 ment methods on non-uniform meshes. *Journal of Computational Physics*  
459 231 (14), 4694–4708.

460 Belytschko, T., Krongauz, Y., Organ, D., Fleming, M., Krysl, P., 1996. Mesh-  
461 less methods: an overview and recent developments. *Computer methods*  
462 *in applied mechanics and engineering* 139 (1), 3–47.

463 Beuchler, S., 2013. Inexact additive Schwarz solvers for hp-fem discretizations  
464 in three dimensions. In: Apel, T., Steinbach, O. (Eds.), *Advanced Finite*  
465 *Element Methods and Applications*. Vol. 66 of *Lecture Notes in Applied*  
466 *and Computational Mechanics*. Springer Berlin Heidelberg, pp. 91–108.

467 Bientinesi, P., Eijkhout, V., Kim, K., Kurtz, J., Van De Geijn, R., 2010a.  
468 Sparse direct factorizations through unassembled hyper-matrices. *Com-*  
469 *puter Methods in Applied Mechanics and Engineering* 199 (9), 430–438.

470 Bientinesi, P., Eijkhout, V., Kim, K., Kurtz, J., van de Geijn, R., 2010b.  
471 Sparse direct factorizations through unassembled hyper-matrices. *Com-*  
472 *puter Methods in Applied Mechanics and Engineering* 199 (9-12), 430 –  
473 438.

- 474 Blackford, S., Dongarra, J., 1999. LAPACK working note 41 – installation  
475 guide for LAPACK. Tech. rep., University of Tennessee.  
476 URL "<http://www.netlib.org/lapack/lawnspdf/lawn41.pdf>"
- 477 Calo, V. M., Collier, N. O., Pardo, D., Paszyński, M. R., 2011. Computa-  
478 tional complexity and memory usage for multi-frontal direct solvers used  
479 in p finite element analysis. *Procedia Computer Science* 4, 1854–1861.
- 480 Collier, N., Dalcin, L., Calo, V., 2014. On the computational efficiency of  
481 isogeometric methods for smooth elliptic problems using direct solvers.  
482 *International Journal for Numerical Methods in Engineering* 100 (8), 620–  
483 632.
- 484 Collier, N., Dalcin, L., Pardo, D., Calo, V., 2013. The cost of continuity:  
485 performance of iterative solvers on isogeometric finite elements. *SIAM J.*  
486 *Scientific Computing* 35 (2), A767 – A784.
- 487 Collier, N., Pardo, D., Dalcin, L., Paszynski, M., Calo, V., 2012. The cost  
488 of continuity: A study of the performance of isogeometric finite elements  
489 using direct solvers. *Computer Methods in Applied Mechanics and Engi-*  
490 *neering* 213-216 (0), 353 – 361.
- 491 Collier, N., Simkins Jr, D. C., 2009. The quasi-uniformity condition for repro-  
492 ducing kernel element method meshes. *Computational Mechanics* 44 (3),  
493 333–342.
- 494 Gao, L., 2013. Kronecker products on preconditioning. Ph.D. thesis, King  
495 Abdullah University of Sciences and Technology (KAUST).
- 496 Guyan, R. J., 1965. Reduction of stiffness and mass matrices. *AIAA journal*  
497 3 (2), 380–380.
- 498 Hughes, T. J., 2012. *The finite element method: linear static and dynamic*  
499 *finite element analysis*. Courier Corporation.
- 500 Hughes, T. J., Cottrell, J. A., Bazilevs, Y., 2005. Isogeometric analysis: CAD,  
501 finite elements, nurbs, exact geometry and mesh refinement. *Computer*  
502 *methods in applied mechanics and engineering* 194 (39), 4135–4195.
- 503 Kirby, R. M., Sherwin, S. J., Cockburn, B., 2012. To cg or to hdg: a com-  
504 parative study. *Journal of Scientific Computing* 51 (1), 183–212.

- 505 Melenk, J. M., Gerdes, K., Schwab, C., 2001. Fully discrete hp-finite el-  
506 ements: fast quadrature. *Computer Methods in Applied Mechanics and*  
507 *Engineering* 190 (32), 4339–4364.
- 508 Nguyen, N., Peraire, J., Cockburn, B., 2011. Hybridizable discontinuous  
509 galerkin methods. In: *Spectral and High Order Methods for Partial Dif-*  
510 *ferential Equations*. Springer, pp. 63–84.
- 511 Orszag, S. A., 1980. Spectral methods for problems in complex geometries.  
512 *Journal of Computational Physics* 37 (1), 70–92.
- 513 Pardo, D., April 2004. Integration of *hp*-adaptivity with a two grid solver:  
514 applications to electromagnetics. Ph.D. thesis, The University of Texas at  
515 Austin.
- 516 Pardo, D., Calo, V., Torres-Verdin, C., Nam, M., 2008a. Fourier series ex-  
517 pansion in a non-orthogonal system of coordinates for the simulation of  
518 3D-DC borehole resistivity measurements. *Computer Methods in Applied*  
519 *Mechanics and Engineering* 197 (21), 1906–1925.
- 520 Pardo, D., Demkowicz, L., 2006. Integration of *hp*-adaptivity with a two grid  
521 solver for elliptic problems. *Computer Methods in Applied Mechanics and*  
522 *Engineering* 195, 674–710.
- 523 Pardo, D., Demkowicz, L., Gopalakrishnan, J., 2006. Integration of *hp*-  
524 adaptivity and a two grid solver for electromagnetic problems. *Computer*  
525 *Methods in Applied Mechanics and Engineering* 195, 2533–2573.
- 526 Pardo, D., Torres-Verdín, C., Nam, M., Paszynski, M., Calo, V., 2008b.  
527 Fourier series expansion in a non-orthogonal system of coordinates for the  
528 simulation of 3D alternating current borehole resistivity measurements.  
529 *Computer Methods in Applied Mechanics and Engineering* 197 (45), 3836–  
530 3849.
- 531 Schoberl, J., Melenk, J. M., Pechstein, C., Zaglmayr, S., 2008. Additive  
532 Schwarz preconditioning for p-version triangular and tetrahedral finite el-  
533 ements. *IMA Journal of Numerical Analysis* 28 (1), 1–24.
- 534 Schoberl, J., Melenk, J. M., Pechstein, C. G. A., Zaglmayr, S. C., 2007.  
535 Schwarz preconditioning for high order simplicial finite elements. In: Wid-  
536 lund, O. B., Keyes, D. E. (Eds.), *Domain Decomposition Methods in Sci-*

- 537      ence and Engineering XVI. Vol. 55 of Lecture Notes in Computational  
538      Science and Engineering. Springer Berlin Heidelberg, pp. 139–150.
- 539      Sourbier, F., Haidar, A., Giraud, L., Ben-Hadj-Ali, H., Operto, S., Virieux,  
540      J., 2011. Three-dimensional parallel frequency-domain visco-acoustic wave  
541      modelling based on a hybrid direct/iterative solver. *Geophysical Prospect-*  
542      *ing* 59 (5), 834–856.
- 543      Vejchodsky, T., Solin, P., 2008. Static condensation, partial orthogonaliza-  
544      tion of basis functions, and ilu preconditioning in the hp-fem. *Journal of*  
545      *Computational and Applied Mathematics* 218 (1), 192 – 200.
- 546      Vos, P. E., Sherwin, S. J., Kirby, R. M., 2010. From h to p efficiently: Im-  
547      plementing finite and spectral hp element methods to achieve optimal per-  
548      formance for low-and high-order discretisations. *Journal of Computational*  
549      *Physics* 229 (13), 5161–5181.
- 550      Wilson, E. L., 1974. The static condensation algorithm. *International Journal*  
551      *for Numerical Methods in Engineering* 8 (1), 198–203.