Journal of Computational Physics 229 (2010) 5161–5181

Contents lists available at ScienceDirect

Journal of Computational Physics

journal homepage: www.elsevier.com/locate/jcp

From *h* to *p* efficiently: Implementing finite and spectral/*hp* element methods to achieve optimal performance for low- and high-order discretisations

Peter E.J. Vos [a,b], Spencer J. Sherwin [a,*], Robert M. Kirby [c,1]

[a] Department of Aeronautics, Imperial College London, London, UK
[b] Flemish Institute for Technological Research (Vito), Mol, Belgium
[c] School of Computing, University of Utah, Salt Lake City, UT, USA

ARTICLE INFO

Article history:
Received 24 November 2009
Received in revised form 22 March 2010
Accepted 23 March 2010
Available online 28 March 2010

Keywords:
Spectral/*hp* element method
Implementation
Elliptic problems

ABSTRACT

The spectral/*hp* element method can be considered as bridging the gap between the – traditionally low-order – finite element method on one side and spectral methods on the other side. Consequently, a major challenge which arises in implementing the spectral/*hp* element methods is to design algorithms that perform efficiently for both low- and high-order spectral/*hp* discretisations, as well as discretisations in the intermediate regime. In this paper, we explain how the judicious use of different implementation strategies can be employed to achieve high efficiency across a wide range of polynomial orders. Furthermore, based upon this efficient implementation, we analyse which spectral/*hp* discretisation (which specific combination of mesh-size *h* and polynomial order *P*) minimises the computational cost to solve an elliptic problem up to a predefined level of accuracy. We investigate this question for a set of both smooth and non-smooth problems.

1. Introduction

The spectral/*hp* element method combines the geometric flexibility of classical *h*-type finite element or finite volume techniques with the desirable resolution properties of spectral methods. In this approach a polynomial expansion of order *P* is applied to every elemental domain of a coarse finite element type mesh. These techniques have been applied in many fundamental studies of fluid mechanics [1] and more recently have gained greater popularity in the modelling of wave-based phenomena such as computational electromagnetics [2] and shallow water problems [3] – particularly when applied within a Discontinuous Galerkin formulation.

Spectral/*hp* element methods can be considered as a high-order extension of classical – traditionally low-order – finite element methods where convergence is not only possible through reducing the characteristic mesh-size *h* but also through increasing the local polynomial order *P* within an elemental subdomain. However, the concept of high and low-order discretisations can mean very different things to different communities. For example, the seminal works by Zienkiewicz and Taylor [4] and Hughes [5] list examples of elemental expansions only up to third or possibly fourth-order, implying that these

* Corresponding author.
E-mail addresses: peter.vos@imperial.ac.uk (P.E.J. Vos), s.sherwin@imperial.ac.uk (S.J. Sherwin), kirby@cs.utah.edu (R.M. Kirby).
[1] Work accomplished while on sabbatical at Imperial College London.

0021-9991/$ - see front matter © 2010 Elsevier Inc. All rights reserved.
doi:10.1016/j.jcp.2010.03.031

orders are considered to be high-order for the traditional *h*-type finite element community. In contrast the text books of the spectral/*hp* element community [6–9] typically show examples of problems ranging from a low-order bound of minimally fourth-order up to anything ranging from 10th to 15th order polynomial expansions. On the other end of the spectrum, practitioners of global spectral methods [10] would probably consider a 16th order global expansion to be relatively low-order approximation.

One could wonder whether these different definitions of low- and high-order is just inherent to the tradition and lore of each of the communities or whether there are more practical reasons for this distinct interpretation. Proponents of lower order methods might highlight that some problems of practical interest are so geometrically complex that one cannot computationally afford to use high-order techniques on the massive meshes required to capture the geometry. Alternatively, proponents of high-order methods highlight that if the problem of interest can be captured on a computational domain at reasonable cost then using high-order approximations for sufficiently *smooth* solutions will provide a lower accuracy for a given computational cost. If one however probes even further it also becomes evident that the different communities choose to implement their algorithms in a different manner. For example the standard *h*-type finite element community will typically uses techniques such as sparse matrix storage formats (where only the non-zero entries of a global matrix are stored) to represent a global operator. In contrast the spectral/*hp* element community acknowledges that for higher polynomial expansions more closely coupled computational work takes place at the individual elemental level and this leads to the use of elemental operators rather than global matrix operators. In addition the global spectral method community often make use of the tensor-product approximations where products of one-dimensional rules for integration and differentiation can be applied. From the results in this paper, it will appear that each of the different implementation strategies will perform poorly when applied outside the aforementioned polynomial regimes typically adopted by the different communities, hinting that the traditional views of low- and high-order may be have been strengthened by these practical barriers.

In this paper, we are therefore lead to ask when we should adopt these different implementation strategies if we are to allow the order of our spatial approximations to vary from $P = 1$ up to say $P = 15$? We note that analytic estimates of computational work in this polynomial regime are difficult if not impossible to establish since the computational effort is highly dependent on hard to predict hardware characteristics such as memory management and caching effect as well as optimised linear algebra packages such as BLAS and LAPACK. We therefore will mainly follow a computational approach to assess the efficiency of the different implementation strategies. The support of various implementation strategies within a spectral/*hp* code will allow the user to cross the community dependent barriers of low- and high-order in an efficient way such as the aforementioned example of $P = 4$ (the high-order limit for traditional finite elements and the low-order limit for spectral/*hp* elements). This surrounding polynomial regime ($2 < P < 6$) is however potentially an optimal/desirable range for applications where the mesh resolution is such that increasing polynomial order leads to the onset of rapid/spectral convergence. This level of resolution might be necessary to capture, for example, a complex geometry. The benefit of intermediate polynomial resolution will however only be observed if one can efficiently implement these polynomial discretisations.

Finally, we can also question whether it is sufficient to know which implementation strategy is optimal in terms of CPU time for a specific polynomial order discretisation? Probably a more pertinent question is consider *given the most efficient implementation, what is the best spectral/hp discretisation to obtain a fixed error for a minimal computational cost?* Since computational cost is impacted by different discretisation methodologies such as element size *h*, polynomial order *P*, adaptive refinement *r* [11] or even the continuity of the approximation *k* [12], there are clearly many factors to consider. To help reduce this extensive parameter space in this paper we will restrict ourselves to just *h* and *P* refinements leading to the question: which specific combination of mesh-size *h* and polynomial order *P* requires the minimal computational cost (*i.e.* run-time) to solve a problem up to a predefined accuracy? We will investigate this question for a set of both smooth and non-smooth elliptic problems. To outline the scope of this study, we must also consider which part of the implementation we look to optimise. In the solution of time-dependent partial differential equations, such as those that arise in fluid mechanics, electromagnetics and oceanography, it is often the case that you can have the repeated application of a matrix problem. However, in a steady partial differential equation which might arise in structural mechanics the cost of setting up and the matrix problem may be as equally important as the solution time. In the following analysis we will adopt the first case and assume the repeated application of the matrix operators is the dominant cost. As a result, we aim to optimise the evaluation of such matrix operators for minimal run-time, thereby neglecting any matrix construction time. In addition, we do not consider any memory constraints.

The paper is organised as follows: in Section 2, we begin by introducing the spectral/*hp* element method and highlighting some of its aspects relevant to the topic of this paper. In Section 3, we explain how a local operator can be efficiently evaluated for both low- and high-order expansions. Therefore, we first introduce and discuss three different implementation strategies using global matrices, local matrices of a sum factorisation technique. We then provide theoretical cost estimates for each strategy. Next we investigate the effect of the different strategies on the actual run-time by a set of computational tests and analyse which strategy should be selected depending on the polynomial order. Subsequently, in Section 4 we look for the optimal combination of mesh-size *h* and polynomial order *P* that solves a certain problem up to a predefined level of accuracy in a minimal amount of time for four different test-problems. Finally Section 5 summarises and concludes the presented work.