



SMAI-JCM

SMAI JOURNAL OF
COMPUTATIONAL MATHEMATICS

Transformations for Piola-mapped
elements

FRANCIS R. A. AZNARAN, PATRICK E. FARRELL & ROBERT C. KIRBY

Volume 8 (2022), p. 399-437.

<https://doi.org/10.5802/smai-jcm.91>

© The authors, 2022.



*The SMAI Journal of Computational Mathematics is a member
of the Centre Mersenne for Open Scientific Publishing*

<http://www.centre-mersenne.org/>

Submissions at <https://smai-jcm.centre-mersenne.org/ojs/submission>

e-ISSN: 2426-8399





Transformations for Piola-mapped elements

FRANCIS R. A. AZNARAN¹
PATRICK E. FARRELL²
ROBERT C. KIRBY³

¹ Mathematical Institute, University of Oxford, Oxford, OX2 6GG, UK

E-mail address: aznaran@maths.ox.ac.uk

² Mathematical Institute, University of Oxford, Oxford, OX2 6GG, UK

E-mail address: patrick.farrell@maths.ox.ac.uk

³ Department of Mathematics, Baylor University; Sid Richardson Science Building; 1410
S. 4th St.; Waco, TX 76706.

E-mail address: robert_kirby@baylor.edu.

Abstract. The Arnold–Winther element successfully discretizes the Hellinger–Reissner variational formulation of linear elasticity; its development was one of the key early breakthroughs of the finite element exterior calculus. Despite its great utility, it is not available in standard finite element software, because its degrees of freedom are not preserved under the standard Piola push-forward. In this work we apply the novel transformation theory recently developed by Kirby [*SMAI J. Comput. Math.*, 4:197–224, 2018] to devise the correct map for transforming the basis on a reference cell to a generic physical triangle. This enables the use of the Arnold–Winther elements, both conforming and nonconforming, in the widely-used Firedrake finite element software, composing with its advanced symbolic code generation and geometric multigrid functionality. Similar results also enable the correct transformation of the Mardal–Tai–Winther element for incompressible fluid flow. We present numerical results for both elements, verifying the correctness of our theory.

2020 Mathematics Subject Classification. 65N30, 65F08.

Keywords. Finite element method, Piola transform, pullback, linear elasticity, Stokes, reference basis, Firedrake.

1. Introduction

The use of a *reference element* is critical in the evaluation of finite element basis functions and their derivatives. One constructs, by some means, the basis functions on a particular fixed element and obtains the basis functions on an arbitrary cell K through a mapping. For scalar-valued function spaces, this is frequently a simple pullback (change of coordinates). However, vector-valued elements discretizing $H(\text{div})$ or $H(\text{curl})$, as well as their tensor-valued counterparts, typically use a different mapping in order to facilitate enforcing appropriate continuity of only normal or tangential components.

The reference element paradigm is typically employed for elements that satisfy a kind of *equivalence*, where the reference element basis functions map directly to the physical element basis functions under the coordinate change. However, for many elements (both classical and modern), this is not the case, as the degrees of freedom (DOFs) are “mixed up” by the pullback in a cell-dependent manner. In [54], Kirby developed a general theory for obtaining the proper basis in the case of affinely mapped cells. This approach gives the correct nodal basis as a linear combination of the mapped reference element basis functions. This linear combination turns out to be sparse, meaning that applying the theory

The first author is supported by the EPSRC Centre for Doctoral Training in Partial Differential Equations: Analysis and Applications [grant number EP/L015811/1] and The MathWorks, Inc. The second author is supported by EPSRC grants EP/V001493/1 and EP/R029423/1. The third author is supported by National Science Foundation grant 1912653.

<https://doi.org/10.5802/smai-jcm.91>

© The authors, 2022

incurs only a small additional cost during the finite element computation, ensuring that the use of these more exotic elements composes neatly with existing high-level automatic code generation software. Kirby and Mitchell [56] generalized the Firedrake [76] code stack to generate and employ this transformation, giving results for Morley, Hermite, Argyris, and Bell elements. The basis on the reference cell for each of these elements is constructed using FIAT [53].

In this paper, we extend the transformation theory presented in [54] to Piola-mapped $H(\text{div})$ elements, both vector- and tensor-valued. Specifically, we focus on the Mardal–Tai–Winther [64] and Arnold–Winther [11, 12] elements. The vector-valued Mardal–Tai–Winther element provides a discretization of $H(\text{div})$ which is nonconforming in H^1 , and works robustly for the scale of operators that interpolate between Stokes and Darcy flow. Thanks to its discrete Korn inequality, it serves as a locking-free elasticity element [65]. It also provides a low-order divergence-free discretization of the incompressible Stokes and Navier–Stokes equations [50]. It has been implemented for a small number of standalone numerical experiments [28], and implemented partially in Diffpack [60, Ch. 4] and SyFi [63, Ch. 15] (a symbolic alternative to FIAT) and envisioned as part of FIAT [63, Ch. 3]. These latter implementations were not fully usable on arbitrary unstructured meshes, due to the lack of transformation theory which we now provide in this work.

The tensor-valued Arnold–Winther elements are used to discretize the Cauchy stress tensor in the stress-displacement formulation of planar linear elasticity. After decades of effort by many researchers, these elements solved a fundamental challenge of numerical elasticity in providing a stable and convergent discretization of the Hellinger–Reissner principle, which moreover enforces symmetry of the stress tensor exactly. Despite being conceived almost two decades ago, the prior implementations of these Arnold–Winther elements have either required the explicit element-by-element construction of the basis [25], or again, have been for one-off numerical experiments, such as for stress reconstruction or the design of error indicators [37, 66, 69, 78]. Again, FIAT has had partial support for this element, but the lack of reference mappings has prevented its use in practice [63, Ch. 3].

We begin in Section 2 with notation used to define geometry and Piola mappings, and summarize some results on how normal and tangential components are transformed. In Section 3, we survey the applications which motivate the use of the Mardal–Tai–Winther element, before describing several other vector-valued elements of interest; analogously in Section 4, we briefly review the role of the Cauchy stress tensor in linear elasticity and its approximation with finite elements, before defining the Arnold–Winther stress elements. We survey the transformation theory of Kirby [54] and show how it applies in the context of Piola elements in Section 5, which is the main contribution of the paper. A discussion of the application of the Arnold–Winther elements to linear elasticity follows in Section 6; in particular, we prove convergence of Nitsche’s method for the imposition of traction boundary conditions. Section 7 offers an interpretation of Piola transformations in terms of the finite element exterior calculus (FEEC), and reviews the role of FEEC in the development of multigrid preconditioners, as discussed in Section 8, in which we apply patch-based multigrid algorithms to precondition the canonical Stokes–Darcy and Hellinger–Reissner systems for the Mardal–Tai–Winther and Arnold–Winther elements respectively. We briefly discuss the implementation within the Firedrake code stack and present some numerical examples for our newly-enabled elements in Section 9.

2. Piola transformation theory

2.1. Notation and preliminaries

Let $\mathbb{M} = \mathbb{R}^{d \times d}$ denote the space of $d \times d$ matrices, and $\mathbb{S} = \mathbb{R}_{\text{sym}}^{d \times d}$ its symmetric subspace. We shall employ the standard notation for the Sobolev space $H^k(\Omega; \mathbb{X})$ (or $L^2(\Omega; \mathbb{X})$ when $k = 0$), with domain $\Omega \subset \mathbb{R}^d$ and codomain \mathbb{X} , and associated norm $\|\cdot\|_{k,\Omega}$ or simply $\|\cdot\|_k$, as well as the standard spaces $H(\text{div})$, $H(\text{div}; \mathbb{M})$, $H(\text{div}; \mathbb{S})$, $H(\text{rot})$, $H(\text{rot}; \mathbb{S})$, $H(\text{curl})$, and $H(\text{curl}; \mathbb{S})$, where the divergence,

rotation, or curl of a tensor field is defined row-wise. Differential operators with subscript h , such as ∇_h , are meant element-wise. The symbol \lesssim denotes inequality up to a constant which may depend on mesh regularity but not mesh spacing h .

Let $\hat{K} \subset \mathbb{R}^2$ be a reference triangle with vertices $\{\hat{\mathbf{x}}_i\}_{i=1}^3$ and let K be a typical element with vertices $\{\mathbf{x}_i\}_{i=1}^3$. We assume the diffeomorphism $F : \hat{K} \rightarrow K$ to be affine, as shown in Fig. 2.1. We number the edges of a triangle by the convention in [79] as follows. Edge i of any triangle excludes the vertex i and is oriented from the lower-numbered vertex to the higher-numbered one. Each edge \mathbf{e} is a column vector running between two vertices of a triangle, and is also used to denote the edge as a set. Its magnitude is the edge length, and we define the unit tangent to edge \mathbf{e} as $\mathbf{t} = \frac{\mathbf{e}}{\|\mathbf{e}\|}$. To each edge, we associate its normal $\mathbf{n} = R\mathbf{t}$, where

$$R = \begin{bmatrix} & 1 \\ -1 & \end{bmatrix} \quad (2.1)$$

is clockwise rotation. Quantities and differential operators with $\hat{}$ are defined analogously for the reference element. Note that the normal to a triangle may point inward or outward according to this convention, and we allow F to reverse orientation, but triangles sharing an edge will automatically agree on the direction (so that local and global direction coincide). This greatly simplifies the required logic for assembly of finite element DOFs involving normal components. This convention can be extended to simplices of any dimension [10, 79].

Given any function $\hat{\phi}$ defined on \hat{K} , we define its *pullback* by

$$\phi := F^{-*}(\hat{\phi}) := \hat{\phi} \circ F^{-1}. \quad (2.2)$$

The pullback can be used to evaluate ϕ at a point $\mathbf{x} = F(\hat{\mathbf{x}})$ by

$$\phi(\mathbf{x}) = \hat{\phi}(\hat{\mathbf{x}}). \quad (2.3)$$

For scalar-valued spaces, the pullback provides an isomorphism between $H^m(\hat{K})$ and $H^m(K)$ for $m \in \mathbb{Z}_{\geq 0}$; its use is very natural in the context of affine equivalent finite elements such as Lagrange or Crouzeix–Raviart. One can define its basis functions on a fixed reference element \hat{K} and obtain the basis functions on any K by pullback. This can be adapted, with some technicalities, when affine equivalence fails [54]. Other shape spaces, such as the globally defined Rannacher–Turek quadrilateral element [75, Remarks 2 and 5], do not fall under the reference element paradigm.

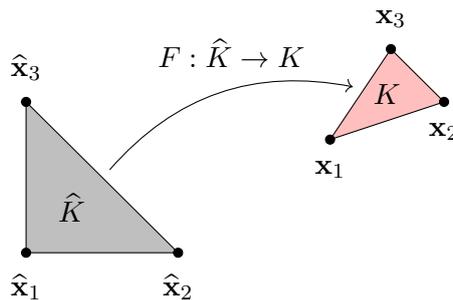


FIGURE 2.1. Affine mapping from a reference cell \hat{K} to a typical cell K .

When working with finite elements discretizing Sobolev spaces such as $H(\text{div})$, it is frequently helpful to work with the contravariant Piola mapping. Piecewise polynomial vector fields will lie in $H(\text{div})$ if and only if they have continuous normal components across interfaces [17, Prop. 2.1.2]; applying this row-wise gives an analogous statement for $H(\text{div}; \mathbb{M})$ and $H(\text{div}; \mathbb{S})$. Hence, finite elements for

$H(\text{div})$ -based spaces tend to include normal components as degrees of freedom, and these will be preserved by using the appropriate mapping.

Definition 2.1. Let $\hat{K}, K \subset \mathbb{R}^d$, let $F : \hat{K} \rightarrow K$ be a diffeomorphism, and let $J(\hat{\mathbf{x}}) = \hat{\nabla} F(\hat{\mathbf{x}})$ be its Jacobian. The *contravariant Piola map* takes $\hat{\Phi} : \hat{K} \rightarrow \mathbb{R}^d$ to $\Phi : K \rightarrow \mathbb{R}^d$ defined by

$$\mathcal{F}^{\text{div}}(\hat{\Phi}) := \frac{1}{\det J} J \hat{\Phi} \circ F^{-1}. \quad (2.4)$$

The *(double) contravariant Piola map* takes $\hat{\tau} : \hat{K} \rightarrow \mathbb{S}$ to $\tau : K \rightarrow \mathbb{S}$ by

$$\mathcal{F}^{\text{div,div}}(\hat{\tau}) := \frac{1}{(\det J)^2} J \left(\hat{\tau} \circ F^{-1} \right) J^T. \quad (2.5)$$

The Piola transforms $\mathcal{F}^{\text{div}} : H(\text{div}, \hat{K}) \rightarrow H(\text{div}, K)$, $\mathcal{F}^{\text{div,div}} : H(\text{div}, \hat{K}; \mathbb{S}) \rightarrow H(\text{div}, K; \mathbb{S})$ are well-defined and isomorphisms. Analogous considerations hold for $H(\text{rot})$, $H(\text{rot}; \mathbb{S})$, $H(\text{curl})$, $H(\text{curl}; \mathbb{S})$, tangential traces, and the covariant Piola mappings:

Definition 2.2. The *covariant Piola map* takes $\hat{\Phi} : \hat{K} \rightarrow \mathbb{R}^d$ to $\Phi : K \rightarrow \mathbb{R}^d$ by

$$\mathcal{F}^{\text{curl}}(\hat{\Phi}) := J^{-T} \hat{\Phi} \circ F^{-1}. \quad (2.6)$$

The *(double) covariant Piola map* takes $\hat{\tau} : \hat{K} \rightarrow \mathbb{S}$ to $\tau : K \rightarrow \mathbb{S}$ by

$$\mathcal{F}^{\text{curl,curl}}(\hat{\tau}) := J^{-T} \left(\hat{\tau} \circ F^{-1} \right) J^{-1}. \quad (2.7)$$

Unlike in the usual definition of the Piola map, the absolute value of the determinant is not taken in (2.4); this is explained in [79], which gives a thorough exposition of transforming and assembling vector finite elements in the case that the spaces and degrees of freedom are exactly preserved under the Piola mappings. The central contribution of this paper is to adapt the techniques of [54] to handle $H(\text{div})$ elements that are *not* exactly preserved under contravariant Piola mapping. This occurs, for example, when elements use degrees of freedom involving edge tangents or vertex values in addition to edge/face normals and internal ones. This enables us to develop general and robust software implementations of these elements.

2.2. Transformation of components and their moments

Now, we review the relationship between the reference and physical element normals and tangents when F is affine. Suppose edge $\mathbf{e} = \mathbf{x}_i - \mathbf{x}_j$ connects vertices i and j of the physical element. Since $F(\hat{\mathbf{x}}_i) = \mathbf{x}_i$, we have

$$\mathbf{e} = \mathbf{x}_i - \mathbf{x}_j = J(\hat{\mathbf{x}}_i - \hat{\mathbf{x}}_j) = J\hat{\mathbf{e}}, \quad (2.8)$$

and equivalently,

$$\|\mathbf{e}\|\mathbf{t} = \|\hat{\mathbf{e}}\|J\hat{\mathbf{t}}. \quad (2.9)$$

In a similar way, we use the identity $\frac{1}{\det J} R J R^T = J^{-T}$ for all 2×2 matrices to obtain

$$\|\mathbf{e}\|\mathbf{n} = R\mathbf{e} = R J \hat{\mathbf{e}} = (\det J) J^{-T} R \hat{\mathbf{e}} = (\det J) J^{-T} \|\hat{\mathbf{e}}\|\hat{\mathbf{n}}. \quad (2.10)$$

In other words, (2.9) shows that the Jacobian maps reference tangents onto physical ones (up to scale factor), but a different mapping given in (2.10) is required to obtain the physical normal from the reference one.

These calculations directly inform how the Piola mapping works. Let us consider some $\hat{\Phi}$ defined over \hat{K} and its image under the Piola map $\Phi = \mathcal{F}^{\text{div}}(\hat{\Phi})$. At each point \mathbf{x} on edge $\mathbf{e} \subset \partial K$ the Piola mapping preserves the normal component scaled by edge length, for we have

$$\|\mathbf{e}\|\Phi(\mathbf{x}) \cdot \mathbf{n} = \left(\frac{1}{\det J} J \hat{\Phi}(\hat{\mathbf{x}}) \right)^T \left(\|\hat{\mathbf{e}}\| (\det J) J^{-T} \hat{\mathbf{n}} \right) = \|\hat{\mathbf{e}}\| \hat{\Phi}(\hat{\mathbf{x}}) \cdot \hat{\mathbf{n}}. \quad (2.11)$$

This calculation also specifies how integral moments of normal components against some function $\mu(\mathbf{x}) = \hat{\mu}(F^{-1}(\mathbf{x}))$ along edges behave under Piola mapping. Changing variables from \mathbf{e} to $\hat{\mathbf{e}}$ introduces a Jacobian factor of $\|\mathbf{e}\|/\|\hat{\mathbf{e}}\|$, so that

$$\int_{\mathbf{e}} (\Phi(s) \cdot \mathbf{n}) \mu(s) ds = \int_{\hat{\mathbf{e}}} \frac{\|\hat{\mathbf{e}}\|}{\|\mathbf{e}\|} (\hat{\Phi}(\hat{s}) \cdot \hat{\mathbf{n}}) \hat{\mu}(\hat{s}) \frac{\|\mathbf{e}\|}{\|\hat{\mathbf{e}}\|} d\hat{s} = \int_{\hat{\mathbf{e}}} (\hat{\Phi}(\hat{s}) \cdot \hat{\mathbf{n}}) \hat{\mu}(\hat{s}) d\hat{s}. \quad (2.12)$$

On the other hand, the Piola map fails to preserve the tangential direction:

$$\|\mathbf{e}\| \Phi(\mathbf{x}) \cdot \mathbf{t} = \left(\frac{1}{\det J} J \hat{\Phi}(\hat{\mathbf{x}}) \right)^T (\|\hat{\mathbf{e}}\| J \hat{\mathbf{t}}) = \frac{\|\hat{\mathbf{e}}\|}{\det J} \hat{\Phi}(\hat{\mathbf{x}})^T J^T J \hat{\mathbf{t}}. \quad (2.13)$$

It is useful to further resolve $J^T J \hat{\mathbf{t}}$ and hence $\|\mathbf{e}\| \Phi(\mathbf{x}) \cdot \mathbf{t}$ in terms of the reference normal and tangential components. We define

$$\hat{G} = \begin{bmatrix} \hat{\mathbf{n}} & \hat{\mathbf{t}} \end{bmatrix}, \quad (2.14)$$

which is an orthogonal matrix. Using $\mathbb{I} = \hat{G} \hat{G}^T$ gives

$$J^T J \hat{\mathbf{t}} = \hat{G} \left(\hat{G}^T J^T J \hat{\mathbf{t}} \right) = \tilde{\alpha} \hat{\mathbf{n}} + \tilde{\beta} \hat{\mathbf{t}}, \quad (2.15)$$

where

$$\begin{bmatrix} \tilde{\alpha} \\ \tilde{\beta} \end{bmatrix} = \hat{G}^T J^T J \hat{\mathbf{t}}. \quad (2.16)$$

Now, with $\alpha = \frac{\tilde{\alpha}}{\det J}$ and $\beta = \frac{\tilde{\beta}}{\det J}$, we have

$$\|\mathbf{e}\| \Phi(\mathbf{x}) \cdot \mathbf{t} = \|\hat{\mathbf{e}}\| \hat{\Phi}(\hat{\mathbf{x}}) \cdot (\alpha \hat{\mathbf{n}} + \beta \hat{\mathbf{t}}). \quad (2.17)$$

So, tangential moments are not preserved under the Piola map, but we have

$$\int_{\mathbf{e}} (\Phi(s) \cdot \mathbf{t}) \mu(s) ds = \int_{\hat{\mathbf{e}}} (\hat{\Phi}(\hat{s}) \cdot (\alpha \hat{\mathbf{n}} + \beta \hat{\mathbf{t}})) \hat{\mu}(\hat{s}) d\hat{s}. \quad (2.18)$$

These properties show that vector-valued elements with boundary degrees of freedom involving only normal components can be directly mapped by the Piola transform. Such elements include the Raviart–Thomas [77], Brezzi–Douglas–Marini [20], and Brezzi–Douglas–Fortin–Marini [19] elements. On the other hand, the Mardal–Tai–Winther element includes tangential as well as normal boundary degrees of freedom, and we can expect complications in mapping them.

We can also use the calculations (2.9) and (2.10) to determine what happens to the various components of a tensor under the Piola map. For some symmetric tensor-valued function $\hat{\tau} : \hat{K} \rightarrow \mathbb{S}$ and $\tau = \mathcal{F}^{\text{div,div}}(\hat{\tau})$, the double-normal component is preserved in an analogous sense to the vector case:

$$\|\mathbf{e}\|^2 \mathbf{n}^T \tau(\mathbf{x}) \mathbf{n} = \left(\|\hat{\mathbf{e}}\| (\det J) J^{-T} \hat{\mathbf{n}} \right)^T \left(\frac{1}{(\det J)^2} J \hat{\tau}(\hat{\mathbf{x}}) J^T \right) \left(\|\hat{\mathbf{e}}\| (\det J) J^{-T} \hat{\mathbf{n}} \right) = \|\hat{\mathbf{e}}\|^2 \hat{\mathbf{n}}^T \hat{\tau}(\hat{\mathbf{x}}) \hat{\mathbf{n}}. \quad (2.19)$$

On the other hand, the normal/tangential and double tangential components are not preserved, as direct calculations verify:

$$\|\mathbf{e}\|^2 \mathbf{n}^T \tau(\mathbf{x}) \mathbf{t} = \frac{\|\hat{\mathbf{e}}\|^2}{\det J} \hat{\mathbf{n}}^T \hat{\tau}(\hat{\mathbf{x}}) (J^T J) \hat{\mathbf{t}}, \quad (2.20)$$

and

$$\|\mathbf{e}\|^2 \mathbf{t}^T \tau(\mathbf{x}) \mathbf{t} = \frac{\|\hat{\mathbf{e}}\|^2}{(\det J)^2} \hat{\mathbf{t}}^T (J^T J) \hat{\tau}(\hat{\mathbf{x}}) (J^T J) \hat{\mathbf{t}}. \quad (2.21)$$

The tangential-normal component $\|\mathbf{e}\|^2 \mathbf{t}^T \tau(\mathbf{x}) \mathbf{n}$ follows from (2.20) by symmetry. Moreover, we can keep \hat{G} as in (2.14) and further resolve these in terms of the reference normal/tangential directions by

$$\|\mathbf{e}\|^2 \mathbf{n}^T \tau(\mathbf{x}) \mathbf{t} = \|\hat{\mathbf{e}}\|^2 \hat{\mathbf{n}}^T \hat{\tau}(\hat{\mathbf{x}}) (\alpha \hat{\mathbf{n}} + \beta \hat{\mathbf{t}}) = \|\hat{\mathbf{e}}\|^2 (\alpha \hat{\mathbf{n}}^T \hat{\tau}(\hat{\mathbf{x}}) \hat{\mathbf{n}} + \beta \hat{\mathbf{n}}^T \hat{\tau}(\hat{\mathbf{x}}) \hat{\mathbf{t}}), \quad (2.22)$$

and

$$\|\mathbf{e}\|^2 \mathbf{t}^T \tau(\mathbf{x}) \mathbf{t} = \|\hat{\mathbf{e}}\|^2 \left(\alpha^2 \hat{\mathbf{n}}^T \hat{\tau}(\hat{\mathbf{x}}) \hat{\mathbf{n}} + 2\alpha\beta \hat{\mathbf{n}}^T \hat{\tau}(\hat{\mathbf{x}}) \hat{\mathbf{t}} + \beta^2 \hat{\mathbf{t}}^T \hat{\tau}(\hat{\mathbf{x}}) \hat{\mathbf{t}} \right). \quad (2.23)$$

Remark 2.3. The tangential and tangential-tangential traces are well-defined at least on polynomial subspaces, and are included here for completeness.

Like in the vector-valued case, we can also show how moments transform. Suppose $\hat{\tau} \in L^2(\hat{K}; \mathbb{S})$ with $\tau = \mathcal{F}^{\text{div}, \text{div}}(\hat{\tau})$ and $\hat{\mu} \in L^2(\hat{\mathbf{e}})$. Integrating both sides of (2.19) and accounting for the Jacobian $ds = \frac{\|\mathbf{e}\|}{\|\hat{\mathbf{e}}\|} d\hat{s}$ gives

$$\|\mathbf{e}\| \int_{\mathbf{e}} \left(\mathbf{n}^T \tau(s) \mathbf{n} \right) \mu(s) ds = \|\hat{\mathbf{e}}\| \int_{\hat{\mathbf{e}}} \left(\hat{\mathbf{n}}^T \hat{\tau}(\hat{s}) \hat{\mathbf{n}} \right) \hat{\mu}(\hat{s}) d\hat{s}. \quad (2.24)$$

Similarly,

$$\|\mathbf{e}\| \int_{\mathbf{e}} \left(\mathbf{n}^T \tau(s) \mathbf{t} \right) \mu(s) ds = \|\hat{\mathbf{e}}\| \left(\alpha \int_{\hat{\mathbf{e}}} \left(\hat{\mathbf{n}}^T \hat{\tau}(\hat{s}) \hat{\mathbf{n}} \right) \hat{\mu}(\hat{s}) d\hat{s} + \beta \int_{\hat{\mathbf{e}}} \left(\hat{\mathbf{n}}^T \hat{\tau}(\hat{s}) \hat{\mathbf{t}} \right) \hat{\mu}(\hat{s}) d\hat{s} \right), \quad (2.25)$$

and

$$\begin{aligned} \|\mathbf{e}\| \int_{\mathbf{e}} \left(\mathbf{t}^T \tau(s) \mathbf{t} \right) \mu(s) ds &= \|\hat{\mathbf{e}}\| \left(\alpha^2 \int_{\hat{\mathbf{e}}} \left(\hat{\mathbf{n}}^T \hat{\tau}(\hat{s}) \hat{\mathbf{n}} \right) \hat{\mu}(\hat{s}) d\hat{s} + 2\alpha\beta \int_{\hat{\mathbf{e}}} \left(\hat{\mathbf{n}}^T \hat{\tau}(\hat{s}) \hat{\mathbf{t}} \right) \hat{\mu}(\hat{s}) d\hat{s} \right. \\ &\quad \left. + \beta^2 \int_{\hat{\mathbf{e}}} \left(\hat{\mathbf{t}}^T \hat{\tau}(\hat{s}) \hat{\mathbf{t}} \right) \hat{\mu}(\hat{s}) d\hat{s} \right). \end{aligned} \quad (2.26)$$

It follows that tensor-valued elements with boundary DOFs involving only normal-normal components can be directly mapped by the Piola transform; this is for example true of the Tangential-Displacement Normal-Normal Stress element [81] and the Hellan–Herrmann–Johnson element [39, 40, 51]. Since the Arnold–Winther elements also have boundary DOFs of other kinds, application of the Piola map alone will not result in the correct dual basis on each physical cell.

Remark 2.4. We briefly comment on the covariant case for $H(\text{curl})$ elements. As is well-known, in 2D, $H(\text{rot})$ and $H(\text{div})$ are isomorphic by simple rotation, although the spaces are quite different in 3D; indeed, in dimensions strictly greater than 2, the analogous identification is always impossible for the vector-valued spaces [62, Prop. 4.1]. It follows that the transformation theory for $H(\text{rot})$ and $H(\text{curl})$ elements, at least in 2D, may be similar *mutatis mutandis* to what we present in this paper.

3. Vector-valued elements

3.1. Motivations

The Mardal–Tai–Winther (MTW) element [64] provides a uniformly robust element for both $H(\text{div})$ and (nonconforming) H^1 . It is a single discretization that is stable and accurate for both Stokes and Darcy flow as well as the scale of operators interpolating between them, including systems arising from transient Stokes-like flow with small timesteps. Due to its lack of vertex degrees of freedom, it is also appropriate for the interface conditions arising in coupled Stokes–Darcy flow [52]. In the context of elasticity, it admits a discrete Korn inequality [65] and serves as a locking-free element for primal linear elasticity [58], and for Biot’s consolidation model in poroelasticity [61].

It is especially significant for the purposes of incompressible flow that the MTW space V_h may be paired with the space Q_h of piecewise constants, so that the resulting velocity-pressure pair $V_h \times Q_h$ satisfies $\text{div } V_h \subset Q_h$, a highly desirable property since it implies mass conservation (i.e. exactly divergence-free velocities) and pressure robustness [50].

The MTW element offers a nonconforming discretization of the two-dimensional Stokes complex, a smoother variant of the classical L^2 de Rham complex [84]:

$$0 \longrightarrow H^2(\Omega)/\mathbb{R} \xrightarrow{\text{curl}} H^1(\Omega; \mathbb{R}^2) \xrightarrow{\text{div}} L^2(\Omega) \longrightarrow 0. \quad (3.1)$$

While it does not form a true subcomplex due to its nonconforming nature, this relates the MTW element to elements for other spaces and gives stability and convergence theory for mixed problems; indeed, the original motivation for the construction of the MTW space was that its divergence-free subspace is exactly the curl of a related H^2 -nonconforming element, studied in [70] and which we denote by W_h :

$$0 \longrightarrow W_h/\mathbb{R} \xrightarrow{\text{curl}} V_h \xrightarrow{\text{div}} Q_h \longrightarrow 0. \tag{3.2}$$

Both sequences (3.1), (3.2) are exact whenever Ω is simply connected. Additionally, this perspective guides the construction of robust smoothers for optimal-order multilevel algorithms [32], as we shall discuss in Sections 7 and 8.

3.2. The Mardal–Tai–Winther element

In this Section, we define several other vector-valued $H(\text{div})$ elements, to discuss their transformations: the Raviart–Thomas [77], Brezzi–Douglas–Marini [20], as well as the Mardal–Tai–Winther [64] elements. Schematics of these elements are shown in Figure 3.1.

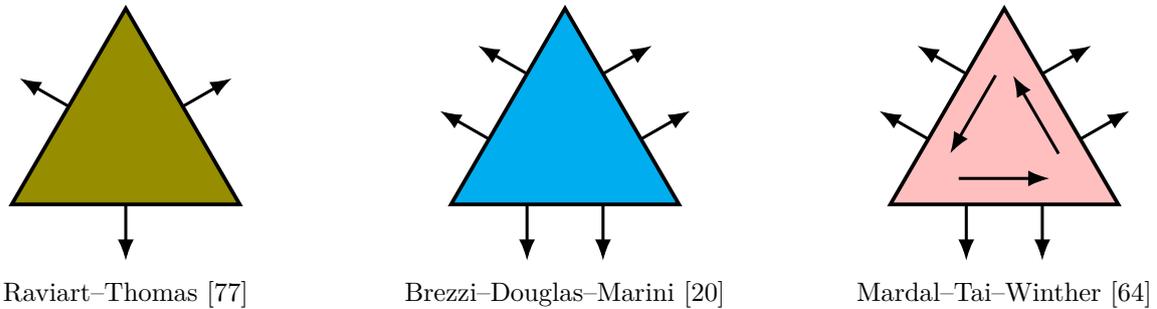


FIGURE 3.1. Some triangular vector $H(\text{div})$ elements. Arrows represent moments of a particular vector component (normal or tangential) along a facet.

Raviart–Thomas elements (the original $H(\text{div})$ element) are based on the function space

$$\text{RT}_m(K) = \mathcal{P}_m(K; \mathbb{R}^2) + \mathbf{x}\mathcal{P}_m(K). \tag{3.3}$$

This is somehow the smallest space of polynomials over K such that the divergence maps exactly onto $\mathcal{P}_m(K)$. The edge degrees of freedom are taken as moments of the normal components up to degree m . The Brezzi–Douglas–Marini elements [20] and Brezzi–Douglas–Fortin–Marini elements [19] have larger spaces than RT, have similar edge degrees of freedom, and so transform in a very similar way. All these elements form equivalent families under the contravariant Piola map.

We let $\{\mu_i\}_{i=0}^m$ be some hierarchically ordered basis for polynomials of degree m on the unit interval (e.g. monomials or Legendre polynomials). (The hierarchical property will not be important for the Raviart–Thomas element, but we shall use it for the Mardal–Tai–Winther space.) We will assume that $\mu_0 = 1$. Let μ_i^k be the mapping of μ_i from the unit interval to edge \mathbf{e}_k of triangle K by standard pullback.

The Raviart–Thomas degrees of freedom include normal moments on each edge:

$$\ell^{\mathbf{n},i,k}(\Phi) = \int_{\mathbf{e}_k} (\Phi(s) \cdot \mathbf{n}) \mu_i^k(s) \, ds \tag{3.4}$$

for each edge \mathbf{e}_k . Now let $\{\hat{\phi}_i\}_{i=1}^{m(m+1)/2}$ be some hierarchically ordered basis for $\mathcal{P}_{m-1}(\hat{K})$ (e.g. the bivariate monomials or Dubiner basis) and $\phi_i = \hat{\phi}_i \circ F^{-1}$ its pullback. We then let $\Phi_{k,i} \in L^2(K; \mathbb{R}^2)$

be the vector-valued function with ϕ_i in component k and 0 in the rest. Defining this, the rest of the degrees of freedom are

$$\ell^{k,i}(\Phi) = \int_K \Phi(\mathbf{x}) \cdot \Phi_{k,i}(\mathbf{x}) \, dx. \tag{3.5}$$

Then, taking the union of all of these gives the nodes for the Raviart–Thomas element. If we construct the nodal basis $\{\hat{\Psi}_i\}$ for $\text{RT}_m(\hat{K})$ or $\text{BDM}_m(\hat{K})$ on a reference element, then a suitable basis for RT_m or BDM_m on any K is obtained by applying \mathcal{F}^{div} to each member of the space. This follows directly from the preservation of normal sense (2.10).

In this context, the Mardal–Tai–Winther element [64] is an example of an $H(\text{div})$ element that includes both normal and tangential degrees of freedom and so requires a more complex reference mapping. On each cell K , the MTW space is

$$\text{MTW}(K) = \left\{ \Phi \in \mathcal{P}_3(K; \mathbb{R}^2) \mid \text{div } \Phi \in \mathcal{P}_0(K), (\Phi \cdot \mathbf{n})|_{\mathbf{e}_k} \in \mathcal{P}_1(\mathbf{e}_k), 1 \leq k \leq 3 \right\}, \tag{3.6}$$

which includes all linear polynomials (and some higher-order terms) and has dimension exactly nine. Six degrees of freedom are the constant and linear moments of the normal components of each edge, as in (3.4). Instead of internal moments like (3.5) for Raviart–Thomas, the remaining three are the constant tangential moment on each edge \mathbf{e}_k :

$$\ell^{\text{t},0,k}(\Phi) = \int_{\mathbf{e}_k} \Phi(\mathbf{x}) \cdot \mathbf{t} \, ds. \tag{3.7}$$

The globally defined MTW space is therefore fully continuous across edges in the normal component, and continuous in a weaker sense in the tangential component, hence is “almost” H^1 -conforming; moreover, application to problems posed in H^1 are simplified in that no facet integral penalty terms are required.

This element is somewhat more expensive than $H(\text{div})$ elements with comparable accuracy. Like BDM_1 , it pairs with piecewise constant pressures but has 50% more degrees of freedom (three per edge rather than two). It gives a second order approximation with only one more degree of freedom than RT_2 , and unlike RT_2 serves also as a H^1 -nonconforming element, but RT_2 pairs with piecewise linear pressures.

4. Tensor-valued elements

4.1. Elastic stress and the numerical enforcement of symmetry

For practical applications of linear elasticity, one is often interested in computing the Cauchy stress tensor with at least as much accuracy as the displacement, and hence, in mixed stress-displacement formulations in which the stress is computed directly rather than, for example, via numerical differentiation after the fact. The mixed dual stress-displacement formulation achieves this, and offers a way to alleviate the well-known numerical phenomenon of volumetric locking in the incompressible limit.

Let $\Omega \subset \mathbb{R}^d$ be a polygonal elastic body; we focus on the planar case $d = 2$. Given Lamé parameters $\lambda, \mu > 0$, the compliance tensor $\mathcal{A} : \mathbb{S} \rightarrow \mathbb{S}$ is, in the homogeneous isotropic case, defined as

$$\mathcal{A}\tau := \frac{1}{2\mu} \left(\tau - \frac{\lambda}{2\mu + d\lambda} (\text{tr } \tau) \mathbb{I} \right). \tag{4.1}$$

Let Γ_D, Γ_N partition $\partial\Omega$, and let $\langle \cdot, \cdot \rangle_{\Gamma_D}$ denote the $(H^{-1/2} \times H_{00}^{1/2})(\Gamma_D; \mathbb{R}^d)$ dual pairing. Given an external body force $f \in L^2(\Omega; \mathbb{R}^d)$, displacement data $u_0 \in H_{00}^{1/2}(\Gamma_D; \mathbb{R}^d)$, and traction data $g \in H^{-1/2}(\Gamma_N; \mathbb{R}^d)$, we consider the Hellinger–Reissner system, which seeks a stress-displacement pair

$(\sigma, u) \in H(\operatorname{div}; \mathbb{S}) \times L^2(\Omega; \mathbb{R}^d)$ satisfying

$$\sigma \mathbf{n} = g \text{ on } \Gamma_N \text{ in the trace sense,}^1 \quad (4.2)$$

and which are critical points of the Hellinger–Reissner functional

$$\mathcal{H}(\sigma, u) := \int_{\Omega} \frac{1}{2} \mathcal{A} \sigma : \sigma + (\operatorname{div} \sigma - f) \cdot u \, dx - \langle \sigma \mathbf{n}, u_0 \rangle_{\Gamma_D}. \quad (4.3)$$

A stationary point (σ, u) satisfies

$$\begin{aligned} \int_{\Omega} \mathcal{A} \sigma : \tau + (\operatorname{div} \tau) \cdot u + (\operatorname{div} \sigma) \cdot v \, dx &= \int_{\Omega} f \cdot v \, dx + \langle \tau \mathbf{n}, u_0 \rangle_{\Gamma_D} \\ \forall (\tau, v) \in H(\operatorname{div}; \mathbb{S}) \times L^2(\Omega; \mathbb{R}^d) \text{ with } \tau \mathbf{n}|_{\Gamma_N} &= 0, \end{aligned} \quad (4.4)$$

i.e. the saddle point system

$$\mathcal{A} \sigma = \varepsilon(u) \quad \text{in } \Omega, \quad (4.5a)$$

$$\operatorname{div} \sigma = f \quad \text{in } \Omega, \quad (4.5b)$$

$$u = u_0 \quad \text{on } \Gamma_D, \quad (4.5c)$$

$$\sigma \mathbf{n} = g \quad \text{on } \Gamma_N, \quad (4.5d)$$

where ε denotes the symmetric gradient (or linearized strain). Here, the constitutive relation (4.5a) is associated with the minimization among stress fields of an appropriate complementary energy, and is formally equivalent to the standard ‘‘Hooke’s law’’ $\sigma = \mathbb{C} \varepsilon(u) = 2\mu \varepsilon(u) + \lambda(\operatorname{tr} \varepsilon(u)) \mathbb{I}$ (where $\mathbb{C} = \mathcal{A}^{-1}$ denotes the elasticity tensor), which in primal formulations is typically used to *define* (and thus eliminate) the stress. However, only the former stress-strain relation (4.5a) remains valid in the incompressible limit $\lambda \gg \mu$.

That the Cauchy stress field σ is symmetric is equivalent to the conservation of angular momentum, and is highly desirable but notoriously difficult to preserve in finite element discretizations; a further challenge arises from the requirement of $H(\operatorname{div}; \mathbb{M})$ -conformity, which ensures that the traction forces on a mesh face shared between two elements are in equilibrium. Before defining the $H(\operatorname{div}; \mathbb{S})$ -discretizing, exactly symmetric Arnold–Winther (AW) elements, we briefly review other approaches to the numerical enforcement of symmetry; for further references, see [17, Ch. 9.3–9.4].

Conventionally, symmetry is enforced weakly via $L^2(\Omega; \mathbb{M})$ -orthogonality to a skew-symmetric subspace; this approach encapsulates the PEERS element [6], and the weakly symmetric successors to the AW elements [9]. Weak imposition of symmetry also fits naturally into the FOSLS (first order system least-squares) method by penalization of asymmetry [22, 68].

A very significant contribution are the TDNNS elements due to Schöberl and Pechstein (née Sinwel) [81]; the stress space consists of the symmetric tensor-valued polynomials of degree $k \geq 1$ whose normal-normal component is continuous across each edge, paired with a displacement space of the same degree whose tangential component is continuous; schematics are provided in Figure 4.1.

¹In this subsection and later from Section 6 onwards, we switch to the usual convention that the unit normal \mathbf{n} to the boundary of Ω or a cell K is outward-pointing, simplifying arguments which employ integration by parts.

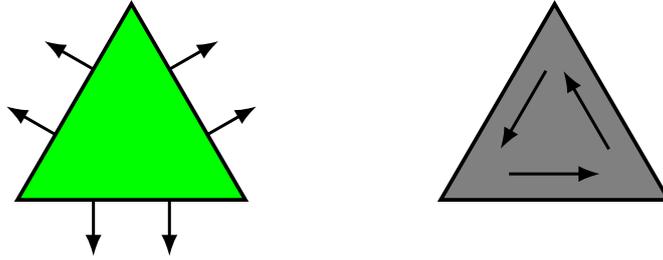


FIGURE 4.1. The Tangential-Displacement and Normal-Normal-Stress mixed element [81]. Thin arrows on the tensor diagram refer to a single tensor component in the given direction (e.g. $\mathbf{n}^T \boldsymbol{\tau} \mathbf{n}$).

Another exactly symmetric element is the Hellan–Herrmann–Johnson (HHJ) element, consisting of symmetric matrix-valued polynomials of degree $k \geq 0$ with continuity of the normal-normal components (that is, $\mathbf{n}^T \boldsymbol{\tau} \mathbf{n}$ is continuous across edges). A schematic is provided in Figure 4.2.

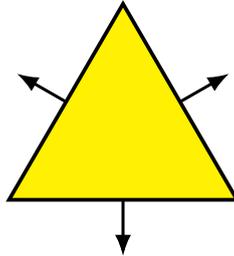


FIGURE 4.2. The Hellan–Herrmann–Johnson element [39, 40, 51].

The TDNNS elements are implemented in Netgen/NGSolve, while the HHJ element was introduced into FEniCS in [38] and later into Firedrake; both implementations were possible without the theory we develop in this paper since both have boundary degrees of freedom involving only the normal-normal components and hence, in light of (2.19), can be Piola-mapped in a standard way.

4.2. The Arnold–Winther elements

Arnold and Winther proposed two exactly symmetric elements discretizing $H(\text{div}; \mathbb{S})$, one conforming [11] and the other nonconforming [12]. We first consider the conforming element. In the lowest-order case, the space is the symmetric matrix-valued quadratic polynomials, augmented by solenoidal cubics:

$$\begin{aligned} \text{AW}^c(K) &= \left\{ \boldsymbol{\tau} \in \mathcal{P}_3(K; \mathbb{S}) \mid \text{div } \boldsymbol{\tau} \in \mathcal{P}_1(K; \mathbb{R}^2) \right\} \\ &= \mathcal{P}_2(K; \mathbb{S}) + \{ \boldsymbol{\tau} \in \mathcal{P}_3(K; \mathbb{S}) \mid \text{div } \boldsymbol{\tau} = 0 \}. \end{aligned} \tag{4.6}$$

This space has dimension 24, and the degrees of freedom are:

- the values of each component of $\boldsymbol{\tau}$ at each vertex of K ,
- the moments of degree 0 and 1 of the normal-normal and normal-tangential components of $\boldsymbol{\tau}$ on each edge,
- the constant moment of each component of $\boldsymbol{\tau}$ over K .

The associated local displacement space is $V_h(K) = \mathcal{P}_1(K; \mathbb{R}^2)$ which is 6-dimensional, with DOFs given by, for example, the values of the 2 components at 3 non-collinear points interior to K . Element diagrams are provided in Figure 4.3.

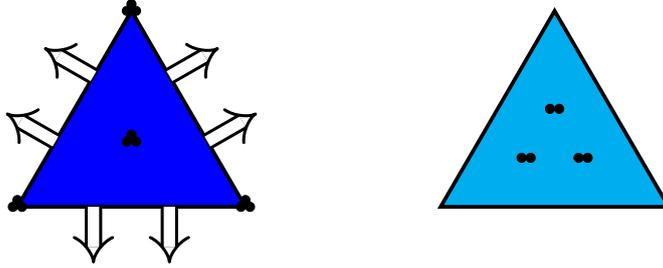


FIGURE 4.3. The conforming Arnold–Winther element [11], and the discontinuous Lagrange element for displacement with which it is paired. Thick arrows refer to both components of a tensor in a given direction (e.g. $\tau \mathbf{n}$ or $\tau \mathbf{t}$). A clutch of circles either indicates internal moments of each unique component or evaluation of each component at a point.

To introduce notation for the stress DOFs, we let

$$\ell^{k,i,j}(\tau) = \tau_{i,j}(\mathbf{x}_k) \quad (4.7)$$

denote evaluation of a particular component of the input τ at one of the vertices of triangle K . Similar to the nodes for vector-valued elements, we define the boundary moments by

$$\ell^{\mathbf{n},\mathbf{s},i,k}(\tau) = \int_{\mathbf{e}_k} (\mathbf{n} \cdot \tau(s)\mathbf{s}) \mu_i(s) ds, \quad (4.8)$$

where $\mathbf{s} \in \{\mathbf{n}, \mathbf{t}\}$; for the HHJ element, we thus have nodes $\{\ell^{\mathbf{n},\mathbf{n},i,k}\}_{i=0}^1$ for each edge \mathbf{e}_k of the triangle.

For the interior nodes, denoting by $\{\phi_i\}_i$ a hierarchically ordered basis over K with $\phi_0 = 1$, we let $\tau_{n,i,j}$ denote the tensor with ϕ_n in the i, j entry and 0 elsewhere – a natural injection into the space of tensor-valued orthogonal polynomials. We let

$$\ell^{n,i,j}(\tau) = \int_K \tau(\mathbf{x}) : \tau_{n,i,j}(\mathbf{x}) dx. \quad (4.9)$$

Thus, DOFs for the AW^c space are given in full by $\{\ell^{k,1,1}, \ell^{k,1,2}, \ell^{k,2,2}\}_{k=1}^3, \{\ell^{\mathbf{n},\mathbf{n},i,k}, \ell^{\mathbf{n},\mathbf{t},i,k}\}_{i=0,k=1}^1, \{\ell^{0,1,1}, \ell^{0,1,2}, \ell^{0,2,2}\}$.

The nonconforming stress element introduced in [12] avoids the somewhat unusual feature of vertex degrees of freedom and gives a cheaper (though slightly less accurate) method based on the space

$$\text{AW}^{\text{nc}}(K) = \{\tau \in \mathcal{P}_2(K; \mathbb{S}) \mid \mathbf{n}_k \cdot \tau \mathbf{n}_k \in \mathcal{P}_1(\mathbf{e}_k), \quad 1 \leq k \leq 3\}. \quad (4.10)$$

This space, consisting of symmetric tensors of quadratic polynomials subject to a degree reduction in the normal-normal component on each edge, is 15-dimensional, and is determined by $\{\ell^{\mathbf{n},\mathbf{n},i,k}, \ell^{\mathbf{n},\mathbf{t},i,k}\}_{i=0,k=1}^1, \{\ell^{0,1,1}, \ell^{0,1,2}, \ell^{0,2,2}\}$, i.e.

- the constant and linear moments of both the normal-normal and normal-tangential components on each edge,
- the constant moments of the three unique components on the interior.

It is again paired with the displacement space $\mathcal{P}_1(K; \mathbb{R}^2)$; element diagrams are provided in Figure 4.4.

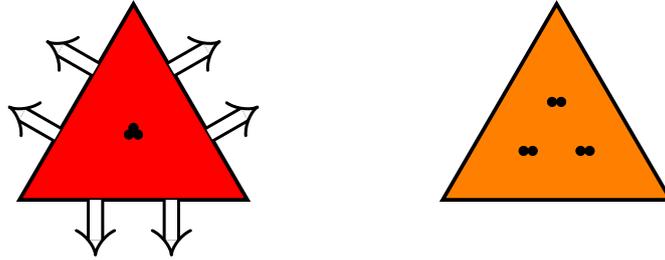


FIGURE 4.4. The nonconforming Arnold–Winther mixed element [12].

The Arnold–Winther elements were one of the first products of the then embryonic finite element exterior calculus [8]. Their novelty is that they form an exact subcomplex of the stress elasticity complex in 2D, with commuting cochain projections:

$$\begin{array}{ccccccccc}
 0 & \longrightarrow & \mathcal{P}_1(\Omega) & \xrightarrow{\subset} & H^2(\Omega) & \xrightarrow{\text{airy}} & H(\text{div}; \mathbb{S}) & \xrightarrow{\text{div}} & L^2(\Omega; \mathbb{R}^2) & \longrightarrow & 0 \\
 & & \downarrow \text{id} & & \downarrow I_h & & \downarrow \Pi_h & & \downarrow P_h & & \\
 0 & \longrightarrow & \mathcal{P}_1(\Omega) & \xrightarrow{\subset} & Q_h & \xrightarrow{\text{airy}_h} & \Sigma_h & \xrightarrow{\text{div}_h} & V_h & \longrightarrow & 0.
 \end{array} \tag{4.11}$$

Here, if (Σ_h, V_h) are either the conforming or nonconforming AW elements, Q_h is either the Argyris space (another element requiring nonstandard transformations) or a certain nonconforming approximation of $H^2(\Omega)$ respectively, airy denotes the Airy stress operator which we will define in (7.12), I_h, Π_h are appropriate densely defined interpolants, and P_h is the L^2 -projection, then it is proved in [11, 12] (see also [85]) that the above diagram is commuting, with exact rows whenever Ω is simply connected. It is these structure-preserving properties which are used to prove unisolvency of dual bases, prove error estimates, and clarify links between their elements and H^2 -(non)conforming elements. Indeed, under the unifying framework, the elasticity system (4.5) is the mixed Hodge Laplacian boundary value problem associated with the load space $L^2(\Omega; \mathbb{R}^2)$ in (4.11), where the base stress space $L^2(\Omega; \mathbb{S})$ is endowed with the energy norm induced by the compliance tensor \mathcal{A} [3, p. 107].

Remark 4.1. Many methods proposed after the seminal papers of Arnold and Winther [11, 12] were, in essence, attempts to construct improved versions of the Arnold–Winther elements. Among these, perhaps the most efficient are the Hu–Zhang elements [46, 48, 49], which strongly and stably enforce symmetry in any dimension; in the 2D conforming case, they have 18 local DOFs in comparison to AW’s 24. The same authors recently provided a unified analysis of methods which strongly impose symmetry [43], using a generalized 4-field formulation of which both the Hu–Zhang and AW elements, among many others, are special cases. These elements also admit an interpretation in terms of the FEEC [27]. The Hu–Zhang elements employ tangential DOFs, and the Piola transformation theory which they therefore require would be very similar to that for the AW elements which we present here.

5. Transforming Piola-inequivalent elements

5.1. Reviewing the transformation theory

Equations (2.13), (2.18), (2.25), and (2.26) indicate that finite elements involving tangential degrees of freedom will not map nicely under the contravariant Piola transform – the reference element nodal basis will not map to the physical element one. In [54], Kirby developed a theory to address the simpler but analogous situation with standard pullback. The broad strokes of the theory go through essentially unchanged if one replaces the regular pullback $F^{-*}(\hat{\phi}) = \hat{\phi} \circ F^{-1}$ with a Piola pullback; for self-containment, we summarize the key ideas here.

We let $\widehat{\Phi}$ be a (column) vector of elements of a function space \widehat{V} over \widehat{K} , taking values in some space such as \mathbb{R}^d or \mathbb{S} , with $\dim \widehat{V} = n$. Let $F : \widehat{K} \rightarrow K$ be the mapping shown in Figure 2.1 and $\mathcal{F}^* : \widehat{V} \rightarrow V$ some pullback operator (e.g. one of the Piola maps) taking \widehat{V} into a function space V defined over K . We define $\mathcal{F}^*(\widehat{\Phi})$ to be the componentwise application of the pullback:

$$\left(\mathcal{F}^*(\widehat{\Phi})\right)_j := \mathcal{F}^*(\widehat{\Phi}_j). \quad (5.1)$$

Then, if we take $\widehat{\Psi}$ to contain the nodal basis functions of \widehat{V} on the reference element and Ψ to contain the nodal basis of V defined on K , an immediate consequence of equations such as (2.18) or (2.25) is that

$$\mathcal{F}^*(\widehat{\Psi}) \neq \Psi, \quad (5.2)$$

and note that equality still fails to hold up to permutation or scaling. If the pullback operator preserves the function space, providing an isomorphism between the instantiations V and \widehat{V} , then $\mathcal{F}^*(\widehat{\Psi})$ will be a basis – just not the nodal one. Consequently, there must exist a (cell-dependent) matrix M such that

$$\Psi = M\mathcal{F}^*(\widehat{\Psi}). \quad (5.3)$$

In this case, one can simply compute the pullback of the reference basis and multiply by M . If M is quite sparse, this is a small additional cost in the finite element computation. With some effort, the theory can be extended to handle the situation when the pullback does not send \widehat{V} onto V , but none of our examples in the present work requires this.

However, as we see from subsection 2.2, it is natural to consider the effect of the pullback on the DOFs, rather than directly on the nodal basis. To consider instead how the dual spaces are transformed, denote by $N = \{\ell_i\}_i$ and $\widehat{N} = \{\widehat{\ell}_i\}_i$ (row) vectors of functionals on V and \widehat{V} respectively, and define the *dual* or *push-forward* $\mathcal{F}_* : V^* \rightarrow \widehat{V}^*$ associated with \mathcal{F}^* by

$$\mathcal{F}_*(\ell) := \ell \circ \mathcal{F}^*. \quad (5.4)$$

Given a vector of dual functionals N , it induces an evaluation operator on n -vectors of functions $\Phi \in V^n$ via the “outer product” $N\Phi := N \otimes \Phi$, where

$$(N \otimes \Phi)_{i,j} := \ell_i(\Phi_j). \quad (5.5)$$

This allows us to extend the push-forward (5.4) to vectors of functionals componentwise either in a manner analogous to (5.1), or equivalently as

$$\mathcal{F}_*(N) := N\mathcal{F}^* = N \otimes \mathcal{F}^*(\cdot). \quad (5.6)$$

Choosing N, \widehat{N} to be the DOFs, then the Kronecker property by which they are characterized is then simply expressed as

$$N \otimes \Psi = \mathbb{I} = \widehat{N} \otimes \widehat{\Psi}. \quad (5.7)$$

Moreover, this clarifies the transformation of the primal basis: expanding the desired physical basis in terms of the mapped reference basis, we have

$$\Psi_j = \sum_{k=1}^n \theta_k^{(j)} \mathcal{F}^*(\widehat{\Psi}_k), \quad j = 1, \dots, n, \quad (5.8)$$

so it follows from the Kronecker property that

$$Q\Theta = \mathbb{I}, \quad (5.9)$$

where

$$Q = N \otimes \mathcal{F}^*(\widehat{\Psi}) = \left(\ell_i(\mathcal{F}^*(\widehat{\Psi}_k))\right)_{i,k=1}^n \quad (5.10)$$

can be interpreted as a generalized Vandermonde matrix and

$$\Theta = Q^{-1} = \left(\theta_k^{(j)} \right)_{k,j=1}^n \quad (5.11)$$

has the vector of coefficients for Ψ_j with respect to $\{\mathcal{F}^*(\widehat{\Psi}_k)\}_k$ in column j .

Now by (5.2) and unisolvency of N , we have

$$N \otimes \mathcal{F}^*(\widehat{\Psi}) \neq N \otimes \Psi = \mathbb{I}, \quad (5.12)$$

so

$$\mathcal{F}_*(N)(\widehat{\Psi}) = N \otimes \mathcal{F}^*(\widehat{\Psi}) \neq \mathbb{I} = \widehat{N} \otimes \widehat{\Psi} = \widehat{N}\widehat{\Psi}, \quad (5.13)$$

so in particular,

$$\mathcal{F}_*(N) \neq \widehat{N}, \quad (5.14)$$

but as before, there must be an invertible P with

$$\widehat{N} = P\mathcal{F}_*(N). \quad (5.15)$$

An important result of [54, Theorem 3.1] is that

$$M = P^T \quad (5.16)$$

(as matrices of numbers). That is, it is sufficient to relate the physical nodes and their push-forwards. As we will see in our examples, it can be more natural to find P^{-1} and then invert it by hand.

5.2. Piola-mapped vector elements

The general approach of the transformation theory developed in [54] is to build the nodal basis on K out of a linear combination of the mapped reference element basis functions. As explained in the previous subsection, for us the transformation of basis functions works by duality, considering instead the push-forward of nodes, which follow readily from the calculations (2.12) and (2.18).

We will use block notation in defining our vectors of degrees of freedom. Let

$$N^k = \left[\ell^{\mathbf{n},0,k} \quad \ell^{\mathbf{t},0,k} \quad \ell^{\mathbf{n},1,k} \right]^T \quad (5.17)$$

contain the three MTW degrees of freedom associated with edge \mathbf{e}_k , and then N is block vector with three parts containing the nodes for each edge:

$$N = [N^1 \quad N^2 \quad N^3]^T. \quad (5.18)$$

Similarly, we let \widehat{N}^k and \widehat{N} contain the corresponding nodes on the reference element. Now, we need to compute the Piola push-forward \mathcal{F}_* of these nodes in terms of the reference element nodes, which effectively builds P^{-1} for the transformation theory.

Let $\widehat{\Phi}$ be defined over \widehat{K} , $1 \leq k \leq 3$, and $i = 0$ or 1 . Then (2.12) gives

$$\mathcal{F}_*(\ell^{\mathbf{n},i,k})(\widehat{\Phi}) = \ell^{\mathbf{n},i,k}(\mathcal{F}^{\text{div}}(\widehat{\Phi})) = \int_{\mathbf{e}_k} \left(\mathcal{F}^{\text{div}}(\widehat{\Phi}) \cdot \mathbf{n} \right) \mu_i \, ds = \int_{\widehat{\mathbf{e}}_k} \left(\widehat{\Phi} \cdot \widehat{\mathbf{n}} \right) \widehat{\mu}_i \, d\widehat{s} = \widehat{\ell}^{\mathbf{n},i,k}(\widehat{\Phi}). \quad (5.19)$$

So, the normal moment nodes are in fact pushed forward to the corresponding reference element normal moment nodes. For tangential moments, (2.18) can be rewritten as

$$\mathcal{F}_*(\ell^{\mathbf{t},i,k})(\widehat{\Phi}) = \alpha^k \widehat{\ell}^{\mathbf{n},i,k}(\widehat{\Phi}) + \beta^k \widehat{\ell}^{\mathbf{t},i,k}(\widehat{\Phi}), \quad (5.20)$$

where we have labeled the α, β with a superscript k indicating they correspond to edge k . We define

$$W^k = \begin{bmatrix} 1 & & \\ \alpha^k & \beta^k & \\ & & 1 \end{bmatrix}, \quad (5.21)$$

and this gives

$$\mathcal{F}_*(N^k) = W^k \widehat{N}^k. \quad (5.22)$$

Using this calculation for each edge, we have

$$\mathcal{F}_*(N) = \begin{bmatrix} W^1 & & \\ & W^2 & \\ & & W^3 \end{bmatrix} \widehat{N}, \quad (5.23)$$

and the block matrix in this last equation is P^{-1} . It is easy to invert each W^k :

$$(W^k)^{-1} = \begin{bmatrix} 1 & & \\ -\frac{\alpha^k}{\beta^k} & \frac{1}{\beta^k} & \\ & & 1 \end{bmatrix} \quad (5.24)$$

and hence obtain P by

$$P = \begin{bmatrix} (W^1)^{-1} & & \\ & (W^2)^{-1} & \\ & & (W^3)^{-1} \end{bmatrix}. \quad (5.25)$$

5.3. Piola-mapped tensor elements

Applying the abstract transformation theory to the Arnold–Winther stress elements introduced above follows a similar pattern, using the results in Section 2. We begin with the nonconforming element, and subsequently consider the inclusion of the vertex DOFs for the conforming element. The existing implementations of these elements [23, 25] require a separate construction of the basis functions for each element, by inverting the Vandermonde matrix arising from the Piola pullback of the primal basis, as in equation (5.9).

For the nonconforming element, we define the vector of nodes associated with edge \mathbf{e}_k by

$$N^k = [\ell^{\mathbf{n},\mathbf{n},0,k} \quad \ell^{\mathbf{n},\mathbf{t},0,k} \quad \ell^{\mathbf{n},\mathbf{n},1,k} \quad \ell^{\mathbf{n},\mathbf{t},1,k}]^T. \quad (5.26)$$

That is, we store the normal and tangential moments of order 0, followed by those of order 1. We also collect the three internal degrees of freedom in the vector

$$N^4 = [\ell^{0,1,1} \quad \ell^{0,1,2} \quad \ell^{0,2,2}]^T, \quad (5.27)$$

so that the nonconforming element has degrees of freedom

$$N = [N^1 \quad N^2 \quad N^3 \quad N^4]^T, \quad (5.28)$$

with a similar definition and ordering of reference element degrees of freedom \widehat{N}^k and \widehat{N} .

Much as with the MTW element, we can use (2.24) and (2.25) to write

$$\mathcal{F}_*(N^k) = W^k \widehat{N}^k, \quad (5.29)$$

where we now have

$$W^k = \frac{\|\widehat{\mathbf{e}}_k\|}{\|\mathbf{e}_k\|} \begin{bmatrix} 1 & & & \\ \alpha^k & \beta^k & & \\ & & 1 & \\ & & \alpha^k & \beta^k \end{bmatrix}, \quad (5.30)$$

and we readily invert W^k by

$$(W^k)^{-1} = \frac{\|\mathbf{e}_k\|}{\|\widehat{\mathbf{e}}_k\|} \begin{bmatrix} 1 & & & \\ -\frac{\alpha^k}{\beta^k} & \frac{1}{\beta^k} & & \\ & & 1 & \\ & & -\frac{\alpha^k}{\beta^k} & \frac{1}{\beta^k} \end{bmatrix} \quad (5.31)$$

and have that

$$P = \begin{bmatrix} (W^1)^{-1} & & & \\ & (W^2)^{-1} & & \\ & & (W^3)^{-1} & \\ & & & \end{bmatrix} \quad (5.32)$$

in a manner very analogous to the MTW element.

Let us apply the push-forward to an internal degree of freedom, recalling that $\phi_0 = 1$:

$$\mathcal{F}_*(\ell^{0,i,j})(\widehat{\tau}) = \int_K \mathcal{F}^{\text{div},\text{div}}(\widehat{\tau}) : \tau_{0,i,j} \, dx = \int_{\widehat{K}} \frac{1}{\det J} (J\widehat{\tau}J^T)_{ij} \, d\widehat{x}. \quad (5.33)$$

If we expand out the quantity $J\widehat{\tau}J^T$, we obtain

$$\begin{bmatrix} J_{11}^2 \widehat{\tau}_{11} + 2J_{11}J_{12}\widehat{\tau}_{12} + J_{12}^2 \widehat{\tau}_{22} & J_{11}J_{21}\widehat{\tau}_{11} + (J_{11}J_{22} + J_{12}J_{21})\widehat{\tau}_{12} + J_{12}J_{22}\widehat{\tau}_{22} \\ * & J_{21}^2 \widehat{\tau}_{11} + 2J_{21}J_{22}\widehat{\tau}_{12} + J_{22}^2 \widehat{\tau}_{22} \end{bmatrix}, \quad (5.34)$$

where the asterisk indicates equality due to symmetry. Using this in (5.33) gives that

$$\begin{aligned} \mathcal{F}_*(\ell^{0,1,1}) &= \frac{1}{\det J} (J_{11}^2 \widehat{\ell}^{0,1,1} + 2J_{11}J_{12}\widehat{\ell}^{0,1,2} + J_{12}^2 \widehat{\ell}^{0,2,2}) \\ \mathcal{F}_*(\ell^{0,1,2}) &= \frac{1}{\det J} (J_{11}J_{21}\widehat{\ell}^{0,1,1} + (J_{11}J_{22} + J_{12}J_{21})\widehat{\ell}^{0,1,2} + J_{12}J_{22}\widehat{\ell}^{0,2,2}) \\ \mathcal{F}_*(\ell^{0,2,2}) &= \frac{1}{\det J} (J_{21}^2 \widehat{\ell}^{0,1,1} + 2J_{21}J_{22}\widehat{\ell}^{0,1,2} + J_{22}^2 \widehat{\ell}^{0,2,2}), \end{aligned} \quad (5.35)$$

so that with

$$\widetilde{W} = \begin{bmatrix} J_{11}^2 & 2J_{11}J_{12} & J_{12}^2 \\ J_{11}J_{21} & J_{11}J_{22} + J_{12}J_{21} & J_{12}J_{22} \\ J_{21}^2 & 2J_{21}J_{22} & J_{22}^2 \end{bmatrix} \quad (5.36)$$

and $\widetilde{W} = \frac{1}{\det J} \widetilde{W}$, we write

$$\mathcal{F}_*(N^4) = \widetilde{W} \widehat{N}^4. \quad (5.37)$$

This gives the result

$$P^{-1} = \begin{bmatrix} W^1 & & & \\ & W^2 & & \\ & & W^3 & \\ & & & \widetilde{W} \end{bmatrix} \quad (5.38)$$

which, again, is readily inverted blockwise to find P . As a remark, the \widetilde{W} block is dense, but denoting its dependence on J by $\widetilde{W} = \widetilde{W}_J$, it is then easily checked that $\widetilde{W}_J^{-1} = \widetilde{W}_{J^{-1}}$, i.e. the inverse can be found by reversing the roles of K and \widehat{K} .

We now turn to the full conforming Arnold–Winther element. It has the same edge and internal nodes to the nonconforming element just considered, but also includes vertex values. However, these transform in a very similar fashion to the internal moments just discussed. We collect the pointwise degrees of freedom for each vertex k into a small vector:

$$N^{0,k} = [\ell^{k,1,1} \quad \ell^{k,1,2} \quad \ell^{k,2,2}]^T. \quad (5.39)$$

We also collect the edge degrees of freedom together in the same way as for the nonconforming element:

$$N^{1,k} = \left[\ell^{\mathbf{n},\mathbf{n},0,k} \quad \ell^{\mathbf{n},\mathbf{t},0,k} \quad \ell^{\mathbf{n},\mathbf{n},1,k} \quad \ell^{\mathbf{n},\mathbf{t},1,k} \right]^T. \quad (5.40)$$

Note that we have added extra superscripts of 0 and 1 to indicate the topological dimension associated with the degrees of freedom. We also include the internal degrees of freedom in a separate vector

$$N^2 = [\ell^{0,1,1} \quad \ell^{0,1,2} \quad \ell^{0,2,2}]^T, \quad (5.41)$$

and hence order the degrees of freedom by

$$N = [N^{0,1} \quad N^{0,2} \quad N^{0,3} \quad N^{1,1} \quad N^{1,2} \quad N^{1,3} \quad N^2]^T. \quad (5.42)$$

The edge and internal degrees of freedom are handled in exactly the same way as for the previous elements. The Piola push-forward of the vertex functional gives

$$\mathcal{F}_*(\ell^{k,i,j})(\hat{\tau}) = \ell^{k,i,j}(\mathcal{F}^{\text{div},\text{div}}(\hat{\tau})) = \left(\frac{1}{(\det J)^2} J \hat{\tau}(\hat{\mathbf{x}}_k) J^T \right)_{ij}. \quad (5.43)$$

This suggests that (except for very special geometry), the Piola transform will not send vertex-oriented basis functions on the reference element to their physical counterpart. However, using the expansion of $J \hat{\tau} J^T$ given in (5.34), a very similar calculation as for the internal moments shows that (5.43) gives for each vertex

$$\mathcal{F}_*(N^{0,k}) = \check{W} \hat{N}^{0,k}, \quad (5.44)$$

where now $\check{W} = \frac{1}{(\det J)^2} \widetilde{W}$, for which again $\check{W}_J^{-1} = \check{W}_{J^{-1}}$. Combining all of the push-forwards gives that

$$\mathcal{F}_*(N) = \begin{bmatrix} \check{W} & & & & & & & \\ & \check{W} & & & & & & \\ & & \check{W} & & & & & \\ & & & W^{1,1} & & & & \\ & & & & W^{1,2} & & & \\ & & & & & W^{1,3} & & \\ & & & & & & \widetilde{W} & \end{bmatrix} \hat{N}. \quad (5.45)$$

Again, the block-diagonal structure of the matrix makes it straightforward to invert and find P .

Remark 5.1. On near-degenerate cells, the analogues of blocks W^k such as (5.21), (5.30) may be ill-conditioned, which could affect the accuracy with which basis functions are evaluated.

5.4. Scale-invariance and conditioning

If global basis functions are of disparate size as the mesh is refined, this introduces an additional source of ill-conditioning into finite element systems. This phenomenon occurs for Hermite and certain other affinely-mapped elements [54], and similar conditions apply to our Piola-mapped elements as well. The solution is to globally re-scale degrees of freedom so that basis functions are of comparable size.

We illustrate these issues for the AW^c element. Let h denote the typical diameter of a cell. Consider a basis function dual to a vertex DOF; it is unity at one vertex in one component and zero at other vertices, hence $\mathcal{O}(1)$ over the cell. A basis function dual to an internal DOF integrates to unity over the cell, hence must be of size $\mathcal{O}(h^{-2})$. Between these cases, a normal moment basis function over an edge must have size $\mathcal{O}(h^{-1})$ over that edge. Integrating the inner products which populate the global mass matrix, we obtain some entries (corresponding to pairs of vertex basis functions) of size $\mathcal{O}(h^2)$, others (corresponding to pairs of internal basis functions) of size $\mathcal{O}(h^{-2})$, and others in between. This leads to an $\mathcal{O}(h^{-4})$ condition number of the mass matrix rather than an $\mathcal{O}(1)$ condition number resulting from, say, standard Lagrange polynomials.

This issue can be readily addressed in a *post hoc* fashion by rescaling the global DOFs according to the local mesh size. For example, one can multiply the edge degrees of freedom by h and the internal degrees of freedom by h^2 , so long as cells sharing a DOF agree on the precise value used in the scaling. This choice was made in our implementation, but alternatively, one could include some power of facet measure in the definitions of degrees of freedom for the finite element. Such scale-invariance is implicit in the treatment of AW^c by Carstensen et al. [25], the Morley–Wang–Xu elements, and related H^3 -nonconforming elements due to Wu and Xu [86] and would give a slightly different transformation matrix.

5.5. Preservation of the constraints

Our three spaces of primary interest can be characterized as standard polynomials of certain total degree, subject to linear constraints. For example, shape functions in the local MTW space (3.6) are cubic, but constrained to have linear normal components on cell edges, and by a loss of degree in their divergence. For each such condition, we may consider a linear functional on vectors/symmetric tensors of polynomials of total degree which vanishes precisely on members of the finite element space. Then, the MTW space may be described as the space of vector-valued cubic polynomials on which all of these constraint functionals vanish, i.e. the intersection of their kernels, and similarly for the AW spaces and symmetric-tensor-valued polynomials.

This is described on the reference element in some detail for the conforming AW^c element in [53]. For the AW^{nc} space, for example, if $\tau \in \mathcal{P}_2(K; \mathbb{S})$, then $\mathbf{n} \cdot \tau \mathbf{n} \in \mathcal{P}_1(\mathbf{e})$ for a given edge \mathbf{e} if and only if it is $L^2(\mathbf{e})$ -orthogonal to any polynomial of degree exactly 2. Thus, within $\mathcal{P}_2(K; \mathbb{S})$, $\text{AW}^{\text{nc}}(K)$ has constraint functionals

$$\ell_k^c(\tau) = \int_{\mathbf{e}_k} (\mathbf{n} \cdot \tau(s) \mathbf{n}) \mu_2^k(s) \, ds, \quad 1 \leq k \leq 3, \quad (5.46)$$

where μ_2^k is the quadratic Legendre polynomial on \mathbf{e}_k . Similarly, one can express the divergence constraints for the MTW element by integrating the divergence of a vector of cubics (which is quadratic) against appropriate Dubiner polynomials.

It is important to note that these constraints are preserved under Piola mappings; that is, if some vector or tensor field satisfies the appropriate conditions on the reference element, then its Piola pullback will satisfy the analogous conditions on the physical element. Due to (2.11), reference element vector fields with linear normal components on edges will map to physical element fields with the same property. Similarly, the constraint on the normal-normal component of the nonconforming AW^{nc} space is preserved in light of (2.19). Considering the degree reduction in the divergence of MTW elements, denote by $\Phi = \mathcal{F}^{\text{div}}(\hat{\Phi})$ the vector Piola map applied to some $\hat{\Phi}$ defined on the reference element. By a direct computation we have $\text{div } \Phi = \frac{1}{\det J} \widehat{\text{div}} \hat{\Phi}$. Similarly, the conforming AW^c element is subject to a loss-of-degree constraint on its (now vector-valued) divergence. Denoting $\tau = \mathcal{F}^{\text{div}, \text{div}}(\hat{\tau})$ for some symmetric $\hat{\tau}$ defined on the reference element, then $\text{div } \tau = \frac{1}{(\det J)^2} J \widehat{\text{div}} \hat{\tau}$. It follows that the Piola maps preserve the degree of the divergence in both cases.

In [54], it was argued that the constraint functionals can moreover be used to construct a basis of the finite element space. A new, larger finite element can be introduced consisting of the unconstrained space of polynomials. The original element's degrees of freedom are then augmented with the particular constraint functionals; together, these nodes form a dual basis for the ambient polynomial space, and the basis functions dual to the original nodes in fact form the nodal basis for the constrained finite element space.

All of the constraint functionals we use here are preserved (up to sign) under the Piola push-forward. Hence, the physical element function space is actually constructed by pullback, even if the basis functions are not preserved. This need not be the case – for example, the constraints on the C^1

Bell element are not preserved under affine pullback, requiring a fuller version of the transformation theory [54].

6. Traction conditions in the Hellinger–Reissner principle with Nitsche’s method

The Arnold–Winther discretization of the Hellinger–Reissner problem (4.5), in the pure displacement case $|\Gamma_N| = 0$, for displacement data u_0 , has discrete weak form given by seeking (σ_h, u_h) in one of the AW pairs $\Sigma_h \times V_h$ such that

$$\begin{aligned} \int_{\Omega} \mathcal{A}\sigma_h : \tau_h + (\operatorname{div}_h \tau_h) \cdot u_h \, dx &= \int_{\partial\Omega} (\tau_h \mathbf{n}) \cdot u_0 \, ds \quad \forall \tau_h \in \Sigma_h, \\ \int_{\Omega} (\operatorname{div}_h \sigma_h) \cdot v_h \, dx &= \int_{\Omega} f \cdot v_h \, dx \quad \forall v_h \in V_h. \end{aligned} \quad (6.1)$$

We now turn to the enforcement of traction conditions (4.2) for the Hellinger–Reissner problem, which for the dual formulation (4.4) are essential (i.e. strongly enforced); we consider a mixed boundary condition ($0 < |\Gamma_D|, |\Gamma_N| < |\partial\Omega|$), although the method may be extended to the pure traction case $|\Gamma_D| = 0$ subject to a quotient of the displacement space by the rigid motions and a suitable compatibility condition.

The original AW papers [11, 12] only treated the case of the elastic body clamped everywhere on the boundary, i.e. $|\Gamma_N| = 0$ with homogeneous displacement data $u_0 \equiv 0$. Carstensen et al. [25] identified (in particular, inhomogeneous) traction conditions as a substantial practical difficulty associated with the AW elements, due to delicate interdependence between DOFs at the boundary, which moreover depend on the shape of the boundary at a given boundary vertex; they were able to enforce them using nodal interpolation [24], Lagrange multipliers [25], or elimination of the boundary DOFs by condensation [23]. When the boundary condition is mixed but the traction data g is zero, the relevant trial and test space for the stress is

$$H_{0,\Gamma_N}(\operatorname{div}, \Omega; \mathbb{S}) := \{\tau \in H(\operatorname{div}, \Omega; \mathbb{S}) \mid \langle \tau \mathbf{n}, v \rangle_{\Gamma_N} = 0 \, \forall v \in H_{0,\Gamma_D}^1(\Omega; \mathbb{R}^2)\} \quad (6.2)$$

[17, Remark 2.1.3]. Wang [85, Lemma III.1] constructed an interpolant of Scott–Zhang type which preserves the traction-free condition on Γ_N under an elliptic regularity assumption, and which was used to prove a discrete inf-sup condition for this case, but no details of the discrete enforcement of this condition were offered in Pasciak and Wang’s application of AW^c to the homogeneous pure traction problem in [73].

There is no clear way to enforce the traction condition in either the Arnold–Winther spaces or the weak formulation of Hellinger–Reissner. We therefore advocate a simpler approach, employing the classical Nitsche’s method [71] to weakly enforce the condition. There is little literature on the application of Nitsche’s method to the enforcement of essential boundary conditions in dual mixed problems. Our approach is similar to [21, 57] for the mixed Poisson problem; see also [82].

Given traction data g , we augment the Hellinger–Reissner functional (4.3) over the discrete spaces $\Sigma_h \times V_h$ with a term incorporating the traction condition (to ensure consistency, since we do not impose the condition on the spaces), and a consistent, quadratic penalty term, seeking the critical point $(\sigma_h, u_h) \in \Sigma_h \times V_h$ of

$$\mathcal{H}_{h,\gamma}(\sigma_h, u_h) := \mathcal{H}(\sigma_h, u_h) - \int_{\Gamma_N} \mathbf{r}_h \cdot u_h \, ds + \frac{\gamma}{2} \delta_h(\mathbf{r}_h), \quad (6.3)$$

where h denotes the characteristic mesh size, $\gamma > 0$ is an h -independent penalty parameter, $\mathbf{r}_h = \sigma_h \mathbf{n} - g$ denotes the traction residual, and

$$\delta_h(\mathbf{r}_h) := \frac{1}{h} \int_{\Gamma_N} \|\mathbf{r}_h\|^2 \, ds \quad (6.4)$$

may be interpreted as a least-squares term penalizing deviation from the traction condition.

The exact traction satisfies $\sigma \mathbf{n} = g$ in $H^{-1/2}(\Gamma_N; \mathbb{R}^2)$, but a penalty term δ_h in terms of the dual norm in $H^{-1/2}(\Gamma_N; \mathbb{R}^2)$ would not aid the analysis, nor is it practical to equivalently penalize the Riesz representative of \mathbf{r}_h in $H_{00}^{1/2}(\Gamma_N; \mathbb{R}^2)$, or to work with the linearization of such penalizations. To ensure consistency of the $L^2(\Gamma_N; \mathbb{R}^2)$ -penalization (6.4), we assume full elliptic regularity of the stress field by assuming that the solution (σ, u) to (4.4) satisfies $(\sigma, u) \in H^1(\Omega; \mathbb{S}) \times H^2(\Omega; \mathbb{R}^2)$, and that $g \in L^2(\Gamma_N; \mathbb{R}^2)$. The latter assumption typically holds in practice for the traction data (or some discrete approximation thereof), while the former holds in the homogeneous isotropic case if Ω is a convex polygon and \mathcal{A} is smooth.

Remark 6.1. The exact displacement u satisfies an unmixed primal formulation of linear elasticity for which $u \in H^1(\Omega; \mathbb{R}^2)$, precisely the weak form of the unmixed Hodge Laplace equation associated with the displacement space $H^1(\Omega; \mathbb{R}^2)$ in the strain complex dual to (4.11), and moreover in fact satisfies $\varepsilon(u) \in H(\text{div}; \mathbb{S})$ when viewed as the solution to the strong form. However, there is no gain in regularity for the stress field by passing to an alternative formulation.

Remark 6.2. The natural choice of exponent for h in (6.4) is $+1$ and not -1 , by dimensionalization and since we expect $\|\mathbf{r}_h\|_{0, \Gamma_N}^2$ to converge one order slower than $\|\mathbf{r}_h\|_{-1/2, \Gamma_N}^2$. However, the term with negative exponent is more naturally interpretable as a penalization, and was found to be more effective in preliminary numerical experiments. This also informs the choice of discrete norms (6.7) below.

Convergence will be proved only for the AW^c element, but a computational example will be included also for AW^{nc}.

Let \mathcal{T}_h denote a quasi-uniform triangulation of Ω , E_h the set of its internal edges, and Λ_h the set of its vertices.² Linearizing the augmented functional (6.3) over the AW^c pair, we seek $(\sigma_h, u_h) \in \Sigma_h \times V_h$ satisfying

$$\begin{aligned} a_{h,\gamma}(\sigma_h, \tau_h) + b(\tau_h, u_h) &= \int_{\Gamma_D} \tau_h \mathbf{n} \cdot u_0 \, ds + \frac{\gamma}{h} \int_{\Gamma_N} g \cdot \tau_h \mathbf{n} \, ds \quad \forall \tau_h \in \Sigma_h, \\ b(\sigma_h, v_h) &= \int_{\Omega} f \cdot v_h \, dx - \int_{\Gamma_N} g \cdot v_h \, ds \quad \forall v_h \in V_h, \end{aligned} \tag{6.5}$$

where

$$\begin{aligned} a_{h,\gamma}(\sigma_h, \tau_h) &:= \int_{\Omega} \mathcal{A} \sigma_h : \tau_h \, dx + \frac{\gamma}{h} \int_{\Gamma_N} \sigma_h \mathbf{n} \cdot \tau_h \mathbf{n} \, ds, \\ b(\tau_h, v_h) &:= \int_{\Omega} (\text{div } \tau_h) \cdot v_h \, dx - \int_{\Gamma_N} \tau_h \mathbf{n} \cdot v_h \, ds. \end{aligned} \tag{6.6}$$

Define the $H^1(\Omega; \mathbb{S}) \times H^1(\Omega; \mathbb{R}^2)$ -based mesh-dependent norms

$$\|\|\tau_h\|\|_h^2 := \|\tau_h\|_{0,\Omega}^2 + \frac{1}{h} \|\tau_h \mathbf{n}\|_{0,\Gamma_N}^2, \quad \|v_h\|_h^2 := \|\varepsilon_h(v_h)\|_0^2 + \frac{1}{h} \sum_{\mathbf{e} \in E_h} \|[v_h]\|_{0,\mathbf{e}}^2 + \frac{1}{h} \|v_h\|_{0,\Gamma_D}^2, \tag{6.7}$$

where $[[\cdot]]$ on an interior edge $\mathbf{e} = \partial K \cap \partial K'$ denotes the jump $[[v_h]] := v_h|_K - v_h|_{K'}$, and on an exterior edge $\mathbf{e} \subset \partial\Omega$ denotes the identity. We now prove well-posedness of the augmented discrete formulation (6.5) using the standard Brezzi conditions [17, Section 4.2.3]. We take $\gamma \geq 1$.

Proposition 6.3. *The augmented Nitsche system (6.5) is well-posed, uniformly in h , with respect to the norms (6.7).*

Proof. There exists $C_{\mathcal{A}} > 0$ with $\int_{\Omega} \mathcal{A} \tau : \tau \, dx \leq C_{\mathcal{A}}^2 \|\tau\|_{0,\Omega}^2 \, \forall \tau \in L^2(\Omega; \mathbb{S})$. For $\sigma_h, \tau_h \in \Sigma_h, v_h \in V_h$, we have

$$|a_{h,\gamma}(\sigma_h, \tau_h)| \leq \max\{C_{\mathcal{A}}^2, \gamma\} \|\|\sigma_h\|\|_h \|\|\tau_h\|\|_h, \tag{6.8}$$

²We assume the edges of \mathcal{T}_h align with the partition $\Gamma_D \cup \Gamma_N$, so that each external edge lies entirely in one of these.

and

$$\begin{aligned}
 b(\tau_h, v_h) &= - \int_{\Gamma_N} \tau_h \mathbf{n} \cdot v_h \, ds + \sum_{K \in \mathcal{T}_h} \int_{\Omega} (\operatorname{div} \tau_h) \cdot v_h \, dx \\
 &= \sum_K \left(\int_{\partial K} \tau_h \mathbf{n} \cdot v_h \, ds - \int_K \varepsilon(v_h) : \tau_h \, dx \right) - \int_{\Gamma_N} \tau_h \mathbf{n} \cdot v_h \, ds \\
 &= - \int_{\Omega} \varepsilon_h(v_h) : \tau_h \, dx + \sum_{\mathbf{e} \in E_h} \int_{\mathbf{e}} \tau_h \mathbf{n} \cdot \llbracket v_h \rrbracket \, ds + \int_{\Gamma_D} \tau_h \mathbf{n} \cdot v_h \, ds,
 \end{aligned} \tag{6.9}$$

so

$$|b(\tau_h, v_h)| \leq \|\varepsilon_h(v_h)\|_0 \|\tau_h\|_0 + \sum_{\mathbf{e} \in E_h} \|\tau_h \mathbf{n}\|_{0,\mathbf{e}} \|\llbracket v_h \rrbracket\|_{0,\mathbf{e}} + \|\tau_h \mathbf{n}\|_{0,\Gamma_D} \|v_h\|_{0,\Gamma_D}. \tag{6.10}$$

By the scaling $\|\tau_h \mathbf{n}\|_{0,\mathbf{e}} \lesssim h^{-\frac{1}{2}} \|\tau_h\|_{0,K} \, \forall \mathbf{e} \subset \partial K$, we obtain

$$|b(\tau_h, v_h)| \lesssim \|\tau_h\|_h \|v_h\|_h. \tag{6.11}$$

Clearly, $a_{h,\gamma}$ is $\|\cdot\|_h$ -coercive on all of Σ_h uniformly in h , provided $\gamma \geq 1$. Now fix $0 \neq u_h \in V_h$. Define $\tau_h \in \Sigma_h$ by the AW^c DOFs

$$\tau_h(\mathbf{x}) = 0 \quad \forall \mathbf{x} \in \Lambda_h, \tag{6.12a}$$

$$\int_{\mathbf{e}} \tau_h \mathbf{n} \cdot w_h \, ds = \frac{1}{h} \int_{\mathbf{e}} \llbracket u_h \rrbracket \cdot w_h \, ds \quad \forall \mathbf{e} \in E_h, \forall \mathbf{e} \subset \Gamma_D, w_h \in \mathcal{P}_1(\mathbf{e}; \mathbb{R}^2), \tag{6.12b}$$

$$\int_{\mathbf{e}} \tau_h \mathbf{n} \cdot w_h \, ds = 0 \quad \forall \mathbf{e} \subset \Gamma_N, w_h \in \mathcal{P}_1(\mathbf{e}; \mathbb{R}^2), \tag{6.12c}$$

$$\int_K \tau_h : q_h \, dx = - \int_K \varepsilon(u_h) : q_h \, dx \quad \forall K \in \mathcal{T}_h, q_h \in \mathcal{P}_0(K; \mathbb{S}) = \mathbb{S}. \tag{6.12d}$$

Choosing $w_h = \llbracket u_h \rrbracket \in \mathcal{P}_1(\mathbf{e}; \mathbb{R}^2)$ in (6.12b), $q_h = \varepsilon(u_h) \in \mathbb{S}$ in (6.12d) gives

$$\int_{\mathbf{e}} \tau_h \mathbf{n} \cdot \llbracket u_h \rrbracket \, ds = \frac{1}{h} \|\llbracket u_h \rrbracket\|_{0,\mathbf{e}}^2 \, \forall \mathbf{e} \in E_h, \quad \int_{\mathbf{e}} \tau_h \mathbf{n} \cdot u_h \, ds = \frac{1}{h} \|u_h\|_{0,\mathbf{e}}^2 \, \forall \mathbf{e} \subset \Gamma_D, \tag{6.13}$$

$$\int_K \tau_h : \varepsilon(u_h) \, dx = - \|\varepsilon(u_h)\|_{0,K}^2 \, \forall K \in \mathcal{T}_h. \tag{6.14}$$

Using (6.9),

$$b(\tau_h, u_h) = - \sum_K \left(- \|\varepsilon(u_h)\|_{0,K}^2 \right) + \sum_{\mathbf{e} \in E_h} \frac{1}{h} \|\llbracket u_h \rrbracket\|_{0,\mathbf{e}}^2 + \sum_{\mathbf{e} \subset \Gamma_D} \frac{1}{h} \|u_h\|_{0,\mathbf{e}}^2 = \|u_h\|_h^2. \tag{6.15}$$

It remains to show that $\|u_h\|_h \gtrsim \|\tau_h\|_h$. For every $\mathbf{e} \subset \Gamma_N$, since the DOFs associated to an edge and its endpoints determine $\tau_h \mathbf{n}$ on that edge, we have $\tau_h \mathbf{n} = 0$ by (6.12a), (6.12c), so it suffices to show $\|u_h\|_h \gtrsim \|\tau_h\|_{0,\Omega}$. By (6.12b), (6.12d), we have $\pi_{\mathbf{e}}(\tau_h \mathbf{n}) = \frac{1}{h} \llbracket u_h \rrbracket$, $\pi_K(\tau_h) = -\varepsilon(u_h) \, \forall \mathbf{e} \in E_h, \mathbf{e} \subset \Gamma_D, K \in \mathcal{T}_h$, where $\pi_{\mathbf{e}} : L^2(\mathbf{e}; \mathbb{R}^2) \rightarrow \mathcal{P}_1(\mathbf{e}; \mathbb{R}^2)$, $\pi_K : L^2(K; \mathbb{S}) \rightarrow \mathbb{S}$ are orthogonal projections. By equivalence of norms on $\widehat{\Sigma}_h$ on the reference cell \widehat{K} , we have $\|\widehat{\tau}_h\|_{0,\widehat{K}}^2 \lesssim \sum_{\text{vertices } \widehat{\mathbf{x}} \in \widehat{K}} |\widehat{\tau}_h(\widehat{\mathbf{x}})|^2 + \sum_{\widehat{\mathbf{e}} \subset \partial \widehat{K}} \|\pi_{\widehat{\mathbf{e}}}(\widehat{\tau}_h \widehat{\mathbf{n}})\|_{0,\widehat{\mathbf{e}}}^2 + \|\pi_{\widehat{K}} \widehat{\tau}\|_{0,\widehat{K}}^2$, so by a scaling argument we obtain $\|\tau_h\|_{0,K}^2 \lesssim \sum_{\mathbf{e} \subset \partial K} \frac{1}{h} \|\llbracket u_h \rrbracket\|_{0,\mathbf{e}}^2 + \|\varepsilon(u_h)\|_{0,K}^2$. This shows

$$\inf_{u_h \in V_h} \sup_{\tau_h \in \Sigma_h} \frac{b(\tau_h, u_h)}{\|\tau_h\|_h \|u_h\|_h} \gtrsim 1. \tag{6.16}$$

■

Remark 6.4. In analogy with the RT or BDM spaces considered in [21, 57], we have the useful equilibrium property $\operatorname{div} \Sigma_h \subset V_h$, but in disanalogy we in general have $(\Sigma_h|_{\mathbf{e}}) \mathbf{n} \notin V_h|_{\mathbf{e}}$ on edges \mathbf{e} .

For error estimation, we employ intermediate approximants $\Pi_h \sigma, P_h u$, where $P_h : L^2(\Omega; \mathbb{R}^2) \rightarrow V_h$ denotes the orthogonal projection and $\Pi_h : H^1(\Omega; \mathbb{S}) \rightarrow \Sigma_h$ denotes the Clément-like interpolation operator constructed by Arnold and Winther [11], which enjoy the following approximation properties for all $(\tau, v) \in H^1(\Omega; \mathbb{S}) \times L^2(\Omega; \mathbb{R}^2)$:

- (i) $\operatorname{div} \Pi_h \tau = P_h \operatorname{div} \tau$ (as in a smoothed variant of the commuting diagram (4.11));
- (ii) $(\Pi_h \tau) \mathbf{n} = \pi_{\mathbf{e}}(\tau \mathbf{n})$ on all edges \mathbf{e} ;
- (iii) $\|\tau - \Pi_h \tau\|_{0,\Omega} \lesssim h^m \|\tau\|_{m,\Omega}, 1 \leq m \leq 3$;
- (iv) $\|v - P_h v\|_{0,\Omega} \lesssim h^m \|v\|_{m,\Omega}, 0 \leq m \leq 2$.

Lemma 6.5 (Local approximation in H^1). *For each $K \in \mathcal{T}_h$ and $\tau \in H^2(\Omega; \mathbb{S})$, we have*

$$\|\tau - \Pi_h \tau\|_{1,K} \lesssim h \|\tau\|_{2,S_K}, \quad (6.17)$$

where S_K is a patch of cells neighbouring K such that $\{S_K\}_{K \in \mathcal{T}_h}$ has the finite overlapping property.

Proof. We adapt the proof of (iii) in [11], from which we recall that the error in Π_h may be written as $I - \Pi_h = (I - \Pi_h^0)(I - R_h)$, where $R_h : L^2(\Omega; \mathbb{S}) \rightarrow \Sigma_h \cap H^1(\Omega; \mathbb{S})$ is a Clément interpolant satisfying

$$\|\tau - R_h \tau\|_{j,K} \lesssim h^{m-j} \|\tau\|_{m,S_K}, \quad 0 \leq j \leq 1, \quad j \leq m \leq 3, \quad (6.18)$$

and S_K is a patch of the required form, and $\Pi_h^0 : H^1(\Omega; \mathbb{S}) \rightarrow \Sigma_h$ is the canonical interpolation operator except at the vertices, at which $(\Pi_h^0 \tau)(\mathbf{x}) := 0 \forall \mathbf{x} \in \Lambda_h$. It is easily checked that $\Pi_{\hat{K}}^0$, the restriction of Π_h^0 to a single cell \hat{K} , is bounded from $H^1(\hat{K}; \mathbb{S})$ to $H^r(\hat{K}; \mathbb{S})$ for all $r \geq 0$; choosing $r = 1$, by a scaling argument we obtain $\|\Pi_K^0 \tau\|_{1,K} \lesssim h^{-1} \|\tau\|_{0,K} + \|\tau\|_{1,K}$, so

$$\|\Pi_h^0(I - R_h)\tau\|_{1,K} \lesssim h^{-1} \|(I - R_h)\tau\|_{0,K} + \|(I - R_h)\tau\|_{1,K} \lesssim h \|\tau\|_{2,S_K}, \quad (6.19)$$

which gives the result when combined with (6.18). \blacksquare

We make an additional regularity assumption on the stress field for error analysis.

Proposition 6.6. *Let $(\sigma, u) \in H^2(\Omega; \mathbb{S}) \times H^2(\Omega; \mathbb{R}^2)$. We have the error estimate*

$$\|\sigma - \sigma_h\|_0 \lesssim h(\gamma \|\sigma\|_{2,\Omega} + \|u\|_{2,\Omega}). \quad (6.20)$$

In particular, we expect the traction residual to converge as $\|\sigma_h \mathbf{n} - g\|_{0,\Gamma_N} = \mathcal{O}(h^{\frac{1}{2}})$.

Proof. By the Babuška condition associated with the well-posed system (6.5), applied to $(\sigma_h - \Pi_h \sigma, u_h - P_h u) \in \Sigma_h \times V_h$, there are $(\tau_h, v_h) \in \Sigma_h \times V_h$ with $\|\tau_h\|_h + \|v_h\|_h \leq 1$ and

$$\|\sigma_h - \Pi_h \sigma\|_h + \|u_h - P_h u\|_h \lesssim a_{h,\gamma}(\sigma_h - \Pi_h \sigma, \tau_h) + b(\tau_h, u_h - P_h u) + b(\sigma_h - \Pi_h \sigma, v_h) \quad (6.21)$$

which by consistency of the system (6.5) is equal to

$$a_{h,\gamma}(\sigma - \Pi_h \sigma, \tau_h) + b(\tau_h, u - P_h u) + b(\sigma - \Pi_h \sigma, v_h) \quad (6.22)$$

$$= a_{h,\gamma}(\sigma - \Pi_h \sigma, \tau_h) - \int_{\Gamma_N} \tau_h \mathbf{n} \cdot (u - P_h u) \, ds \quad \text{using (i), (ii)} \quad (6.23)$$

$$= \int_{\Omega} \mathcal{A}(\sigma - \Pi_h \sigma) : \tau_h \, dx + \underbrace{\frac{\gamma}{h} \int_{\Gamma_N} (\sigma - \Pi_h \sigma) \mathbf{n} \cdot \tau_h \mathbf{n} \, ds}_{(*)} - \int_{\Gamma_N} \tau_h \mathbf{n} \cdot (u - P_h u) \, ds = (\dagger). \quad (6.24)$$

By (ii), the term (*) would vanish were it not for Remark 6.4. Employing multiplicative trace inequalities, approximation properties of P_h , and Lemma 6.5, we have

$$(\dagger) \lesssim \|\sigma - \Pi_h \sigma\|_0 \|\tau_h\|_0 + \|\tau_h\|_{0, \Gamma_N} \left(\frac{\gamma}{h} \|(\sigma - \Pi_h \sigma) \mathbf{n}\|_{0, \Gamma_N} + \|u - P_h u\|_{0, \Gamma_N} \right) \quad (6.25)$$

$$\lesssim h^2 \|\sigma\|_{2, \Omega} + h^{\frac{1}{2}} \sum_{\mathbf{e} \in \Gamma_N} \left(\frac{\gamma}{h} \|\sigma - \Pi_h \sigma\|_{0, \mathbf{e}} + \|u - P_h u\|_{0, \mathbf{e}} \right) \quad (6.26)$$

$$\lesssim h^2 \|\sigma\|_{2, \Omega} + h^{\frac{1}{2}} \sum_K \left(\frac{\gamma}{h} \|\sigma - \Pi_h \sigma\|_{0, K}^{\frac{1}{2}} \|\sigma - \Pi_h \sigma\|_{1, K}^{\frac{1}{2}} + \|u - P_h u\|_{0, K}^{\frac{1}{2}} \|u - P_h u\|_{1, K}^{\frac{1}{2}} \right) \quad (6.27)$$

$$\lesssim h^2 \|\sigma\|_{2, \Omega} + h^{\frac{1}{2}} \left(\frac{\gamma}{h} h \|\sigma\|_{1, \Omega}^{\frac{1}{2}} h^{\frac{1}{2}} \|\sigma\|_{2, \Omega} + h \|u\|_{2, \Omega}^{\frac{1}{2}} h^{\frac{1}{2}} \|u\|_{2, \Omega}^{\frac{1}{2}} \right) = h^2 \|\sigma\|_{2, \Omega} + \gamma h \|\sigma\|_{2, \Omega} + h^2 \|u\|_{2, \Omega}, \quad (6.28)$$

so

$$\|\sigma - \sigma_h\|_0 \leq \|\sigma - \Pi_h \sigma\|_0 + \|\sigma_h - \Pi_h \sigma\|_0 \lesssim h^2 \|\sigma\|_2 + \|\sigma_h - \Pi_h \sigma\|_h + \|u_h - P_h u\|_h \quad (6.29)$$

$$\lesssim h^2 \|\sigma\|_2 + \gamma h \|\sigma\|_2 + h^2 \|u\|_2. \quad (6.30)$$

■

7. An exterior calculus perspective

The finite element exterior calculus [8] is a framework for structure-preservation in finite element discretizations, and for constructing finite element spaces as subcomplexes of complexes of function spaces of differential forms. The inventions of the MTW and AW elements we have considered were motivated by the Stokes (3.1) and elasticity (4.11) complexes respectively. In this Section, we consider the application of FEEC to Piola transformation theory, and to the construction of multigrid smoothers.

7.1. Uniform construction of the pullbacks

The Piola transforms (2.4)–(2.7) may be regarded as the analogy of the standard pullback (2.2) for $H(\text{div})$ - and $H(\text{curl})$ -based spaces, but in fact the pullbacks may be defined uniformly in a manner guided by the FEEC [62, p. 35], [3, Section 6.2.5]; we here employ terminology for which we refer the reader to [3, Ch. 6].³

Let us regard the physical and reference cells K, \widehat{K} as submanifolds of dimension d in \mathbb{R}^d , with $F : \widehat{K} \rightarrow K$ a diffeomorphism. Denote by $\text{Alt}^k \mathbb{R}^d$ the space of alternating k -linear forms on \mathbb{R}^d , and for $M \in \{K, \widehat{K}\}$ let $\Lambda^k(M)$ denote the space of differential k -forms on M , of which functions in the Sobolev spaces we consider will be scalar, vector, and tensor proxies, and the operators of vector calculus will be proxies for the exterior derivative $d : \Lambda^k \rightarrow \Lambda^{k+1}$. Scalar fields may be identified with 0-forms or d -forms, and vector fields with 1-forms or $(d-1)$ -forms. We may specify L^2 -integrability of differential form coefficients with $L^2 \Lambda^k(\Omega) := L^2(\Omega; \text{Alt}^k \mathbb{R}^d)$; Sobolev spaces of differential forms may be defined as $H \Lambda^k := \{\omega \in L^2 \Lambda^k \mid d\omega \in L^2 \Lambda^{k+1}\}$, and correspond to the conventional spaces of vector calculus via $H \Lambda^0 \simeq H^1$, $H \Lambda^1 \simeq H(\text{curl})$, $H \Lambda^{d-1} \simeq H(\text{div})$, and $H \Lambda^d \simeq L^2$.⁴

An elastic stress field \mathcal{T} on M may naturally be identified with an $(\text{Alt}^{d-1} \mathbb{R}^d)$ -valued $(d-1)$ -form $\tau \in \Lambda^{d-1}(M; \text{Alt}^{d-1} \mathbb{R}^d)$, which by the proxy of \mathbb{R}^d for $\text{Alt}^{d-1} \mathbb{R}^d$ may be identified with $\Lambda^{d-1}(M; \mathbb{R}^d)$,

³We remark that much of the literature employing Piola transforms does little to motivate their definition, beyond stating the *consequences* of the definitions.

⁴We already employ \mathbb{R}^d as a proxy for the tangent spaces $\{T_{\mathbf{x}} M\}_{\mathbf{x} \in M}$ on which the alternating forms are defined.

since when integrated over a codimension-1 submanifold (such as the boundary of a subdomain), it should give a vector representing force [8][36, p. 618]. Applying the Hodge star gives an element $\star\tau \in \Lambda^1(M; \mathbb{R}^d)$, i.e. a linear map $\mathbb{R}^d \rightarrow \mathbb{R}^d$ (hence, a matrix) at every point of M , which is the classical characterization of stress. Alternatively, \mathcal{T} may be identified with a (symmetric) covariant 2-tensor field in $\Lambda^1(M; \text{Alt}^1 \mathbb{R}^d)$, which to each point assigns a (symmetric) bilinear form on \mathbb{R}^d [62, p. 10].

Given $\hat{\omega} \in \Lambda^k(\hat{K})$, the derivative of the inverse diffeomorphism $J^{-1}(\mathbf{x}) = (F^{-1})'(\mathbf{x})$ at $\mathbf{x} \in K$ induces an element $\omega = J^{-*}\hat{\omega} \in \Lambda^k(K)$ pointwise via $(J^{-*}\hat{\omega})_{\mathbf{x}} := J^{-1}(\mathbf{x})^*\hat{\omega}_{F^{-1}(\mathbf{x})}$, the pullback of $\hat{\omega}$ under F^{-1} , where $*$ denotes the algebraic pullback of a linear map $L : \mathbb{R}^d \rightarrow \mathbb{R}^d$ given by $L^*\eta(v_1, \dots, v_k) := \eta(Lv_1, \dots, Lv_k)$ for $\eta \in \Lambda^k(M)$, $v_i \in \mathbb{R}^d$. For scalar fields representing 0-forms $\hat{\omega} \in H\Lambda^0(\hat{K})$ (i.e. a constant map at each point), the scalar proxy for the resulting 0-form ω is easily seen to be given simply by precomposition with F^{-1} , because $(J^{-*}\hat{\omega})_{\mathbf{x}} = \hat{\omega}_{F^{-1}(\mathbf{x})}$, which for the proxy gives exactly the standard pullback (2.2). For a vector field \hat{w} representing a 1-form $\hat{\omega} \in H\Lambda^1(\hat{K})$, the canonical identification is $\hat{\omega} = \sum_{i=1}^d \hat{w}_i d\hat{\mathbf{x}}^i$, so that $\hat{\omega}_{\hat{\mathbf{x}}}(v) = \hat{w}(\hat{\mathbf{x}}) \cdot v$ (i.e. \hat{w} is the pointwise Riesz representative of $\hat{\omega}$), so

$$(J^{-*}\hat{\omega})_{\mathbf{x}}(v) = \hat{\omega}_{F^{-1}(\mathbf{x})}(J^{-1}(\mathbf{x})v) = \hat{w}(F^{-1}(\mathbf{x})) \cdot (J^{-1}(\mathbf{x})v) = (J^{-T}(\mathbf{x})\hat{w}(F^{-1}(\mathbf{x}))) \cdot v, \quad (7.1)$$

hence the transformed proxy is given by $\mathcal{F}^{\text{curl}}(\hat{w})$. A vector field \hat{w} can alternatively represent a $(d-1)$ -form via

$$\hat{\omega}_{\hat{\mathbf{x}}}(v_1, \dots, v_{d-1}) = \det[\hat{w}(\hat{\mathbf{x}})|v_1| \dots |v_{d-1}|] =: \sum_{i=1}^d \hat{w}_i(\hat{\mathbf{x}})\phi_i(v_1, \dots, v_{d-1}) = \hat{w}(\hat{\mathbf{x}}) \cdot \phi(v_1, \dots, v_{d-1}), \quad (7.2)$$

for which

$$\begin{aligned} (J^{-*}\hat{\omega})_{\mathbf{x}}(v_1, \dots, v_{d-1}) &= \det[\hat{w}(F^{-1}(\mathbf{x}))|J^{-1}v_1| \dots |J^{-1}v_{d-1}|] \\ &= \det(J^{-1}) \det[J\hat{w}(F^{-1}(\mathbf{x}))|v_1| \dots |v_{d-1}|] \\ &= \det[\mathcal{F}^{\text{div}}(\hat{w})(\mathbf{x})|v_1| \dots |v_{d-1}|]. \end{aligned} \quad (7.3)$$

The derivation of the tensor transforms are analogous; a (symmetric) covariant 2-tensor field $\hat{\tau} \in \Lambda^1(\hat{K}; \text{Alt}^1 \mathbb{R}^d)$ has as proxy a (symmetric) matrix field $\hat{\mathcal{T}}$ via the identification $\hat{\tau}_{\hat{\mathbf{x}}}(v_1, v_2) = v_1^T \hat{\mathcal{T}}(\hat{\mathbf{x}})v_2$, so that $(J^{-*}\hat{\tau})_{\mathbf{x}}(v_1, v_2) = v_1^T J^{-T} \hat{\mathcal{T}}(F^{-1}(\mathbf{x})) J^{-1}v_2 = v_1^T \mathcal{F}^{\text{curl, curl}}(\hat{\mathcal{T}})(\mathbf{x})v_2$. A (symmetric) matrix field $\hat{\mathcal{T}}$ may instead serve as a proxy for a (symmetric) form $\hat{\tau} \in \Lambda^{d-1}(\hat{K}; \text{Alt}^{d-1} \mathbb{R}^d)$ via the identification

$$\hat{\tau}_{\hat{\mathbf{x}}}(v_1, \dots, v_{d-1})(z_1, \dots, z_{d-1}) = \phi(v_1, \dots, v_{d-1})^T \hat{\mathcal{T}}(\hat{\mathbf{x}})\phi(z_1, \dots, z_{d-1}), \quad (7.4)$$

where ϕ is defined by (7.2).⁵ By definition of ϕ it is easily checked that $\phi(J^{-1}v_1, \dots, J^{-1}v_{d-1}) = \frac{1}{\det J} J^T \phi(v_1, \dots, v_{d-1})$, and hence

$$\begin{aligned} (J^{-*}\hat{\tau})_{\mathbf{x}}(v_1, \dots, v_{d-1})(z_1, \dots, z_{d-1}) &= \phi(J^{-1}v_1, \dots, J^{-1}v_{d-1})^T \hat{\mathcal{T}}(F^{-1}(\mathbf{x}))\phi(J^{-1}z_1, \dots, z_{d-1}) \\ &= \phi(v_1, \dots, v_{d-1})^T \mathcal{F}^{\text{div, div}}(\hat{\mathcal{T}})(\mathbf{x})\phi(z_1, \dots, z_{d-1}). \end{aligned} \quad (7.5)$$

Note that neither tensor transform is given by the row-wise application of the corresponding vector transform.

We hereafter assume that J is spatially constant, so that the pullback commutes with the exterior derivative. In this case, a crucial property is that moments of the exterior derivatives against appropriate fields are preserved, a primary motivation for the use of the Piola pullbacks in mixed finite

⁵While this work was under review, the preprint [26] was released, which offers a framework for identifying the (differential-form-valued) differential forms to which tensor-valued functions are proxy; our formula (7.4) for $H(\text{div}; \mathbb{S})$ functions coincides (up to signs of terms) with the identification proposed in [26, p. 18].

elements [17, Lemmas 2.1.6, 2.1.9]. Let $\omega \in H\Lambda^k(K)$, $\mu \in H\Lambda^{d-k-1}(K)$ with pullbacks $\widehat{\omega}, \widehat{\mu}$, then since the pullback respects the exterior product and commutes with the exterior derivative,

$$\int_K \mu \wedge d\omega = \int_K J^{-*} \widehat{\mu} \wedge dJ^{-*} \widehat{\omega} = \int_K J^{-*} \widehat{\mu} \wedge J^{-*} d\widehat{\omega} = \int_K J^{-*} (\widehat{\mu} \wedge d\widehat{\omega}) = \pm \int_{\widehat{K}} \widehat{\mu} \wedge d\widehat{\omega}. \quad (7.6)$$

Note that since we have chosen to allow F to reverse orientation, in general the moment preservation (7.6) will only be up to sign. By Stokes' theorem, it also follows that

$$\int_{\partial K} \mu \wedge \omega = \pm \int_{\partial \widehat{K}} \widehat{\mu} \wedge \widehat{\omega}, \quad (7.7)$$

where the boundary integrands are meant in the trace sense.

Combining these identities with the characterization of constrained finite element spaces by kernels of constraint functionals, as in subsection 5.5, also allows for a more elegant verification of their preservation by the Piola maps; the constraints on the divergences of the MTW and AW^c elements are preserved in light of (7.6), while the preservation of boundary constraints follows from the identity (7.7) applied to appropriate $L^2(\partial K)$ -orthogonal polynomials, as constructed e.g. in [72, Section 3].⁶

That the pullback commutes with the exterior derivative also implies the preservation of the kernel of the operators defining the spaces, a property implicit in the classical presentation of the Piola operators (e.g. [67, Section 3.9]). For example, if $K, \widehat{K} \subset \mathbb{R}^2$ (respectively, \mathbb{R}^3), note that if $\widehat{\phi} \in H^1(\widehat{K})$ and $\phi \in H^1(K)$ is its standard pullback given by (2.2), then by the chain rule we have $\nabla \phi = J^{-T} \widehat{\nabla} \widehat{\phi}$. Since $\text{rot} \nabla \equiv 0$ (resp. $\text{curl} \nabla \equiv 0$), certainly both $\widehat{\nabla} \widehat{\phi} \in H(\text{rot}, \widehat{K})$ (resp. $H(\text{curl}, \widehat{K})$) and $\nabla \phi \in H(\text{rot}, K)$ (resp. $H(\text{curl}, K)$), and the covariant map (2.6) “transform[s] vectors of $H(\text{curl}; \Omega)$ like gradients” [17, p. 61], which, on simply connected domains, form the kernel of rot or curl by exactness of an appropriate de Rham sequence; hence

$$\ker(\text{rot}) = \mathcal{F}^{\text{curl}} \left(\ker(\widehat{\text{rot}}) \right) \quad \text{in 2D} \quad \text{and} \quad \ker(\text{curl}) = \mathcal{F}^{\text{curl}} \left(\ker(\widehat{\text{curl}}) \right) \quad \text{in 3D}. \quad (7.8)$$

Analogously, by a direct computation,

$$\text{curl} \phi = \frac{1}{\det J} J \widehat{\text{curl}} \widehat{\phi}, \quad (7.9)$$

which is reflected in (2.4). We now extend the discussion in [67, Section 3.9] to the tensor-valued case in 2D. Let us define the matrix-valued curl of a vector field in 2D,

$$\mathbf{curl} v := \begin{bmatrix} \frac{\partial v_1}{\partial y} & -\frac{\partial v_1}{\partial x} \\ \frac{\partial v_2}{\partial y} & -\frac{\partial v_2}{\partial x} \end{bmatrix}, \quad (7.10)$$

for which

$$\mathbf{curl} v = \frac{1}{\det J} \widehat{\mathbf{curl}} \widehat{v} J^T, \quad (7.11)$$

and the Airy stress function associated with a scalar potential ϕ

$$\text{airy} \phi := \mathbf{curl} \text{curl} \phi = \begin{bmatrix} \frac{\partial^2 \phi}{\partial y^2} & -\frac{\partial^2 \phi}{\partial x \partial y} \\ -\frac{\partial^2 \phi}{\partial x \partial y} & \frac{\partial^2 \phi}{\partial x^2} \end{bmatrix}, \quad (7.12)$$

which is identically symmetric and divergence-free. Then for $\widehat{\phi} \in H^2(\widehat{K})$ with pullback $\phi \in H^2(K)$,

$$\text{airy} \phi = \frac{1}{(\det J)^2} J \left(\widehat{\text{airy}} \widehat{\phi} \right) J^T = \mathcal{F}^{\text{div}, \text{div}} \left(\widehat{\text{airy}} \widehat{\phi} \right), \quad (7.13)$$

⁶The calculations for the divergences of $\widehat{\Phi}, \widehat{\tau}$ in the first half of subsection 5.5 also reduce to the observations that the pullback commutes with the exterior derivative and preserves polynomial degree.

so by exactness of the 2D stress complex (4.11), we obtain that $\ker(\operatorname{div}) = \mathcal{F}^{\operatorname{div}, \operatorname{div}}(\ker(\widehat{\operatorname{div}}))$. Similarly, for the rot of a symmetric matrix field applied row-wise, we calculate that $\nabla^2 \phi = \mathcal{F}^{\operatorname{curl}, \operatorname{curl}}(\widehat{\nabla}^2 \widehat{\phi})$, where $\nabla^2 H^2$ forms the kernel of rot on $H(\operatorname{rot}; \mathbb{S})$ by exactness of the 2D Hessian complex (e.g. [74, Remark 3.16]).

This observation of kernel preservation may also be connected to the (albeit trivial) topologies of K and \widehat{K} : the pullbacks are isomorphisms between appropriate complexes on K and \widehat{K} , which moreover commute with d , hence are cochain maps which preserve the cohomology of each domain.

7.2. Kernel-capturing and robust multigrid

It is now well-established that the characterization of the kernels of discretized differential operators is crucial for the design of robust multigrid schemes [7, 32, 33, 80]; for both the MTW and AW-type elements, this is given precisely by their positions in the discrete exact complexes (3.2) and (4.11), as we now explain. We describe multigrid relaxation in the framework of subspace correction methods [88]. Given a finite-dimensional Hilbert space V of functions on a mesh, consider a decomposition

$$V = \sum_i V_i, \quad (7.14)$$

where the sum need not be direct. The variational problem to be solved over V often takes the form of a symmetric, coercive operator, perturbed by a positive semidefinite singular operator (such as a discretized divergence) which is scaled by some parameter $\alpha > 0$, a physical or penalty parameter which, as it increases, renders the problem difficult to solve. The seminal work of Schöberl [80, Theorem 4.1] revealed sufficient conditions for α -robustness of the parallel subspace correction preconditioner induced by the decomposition (7.14); a key insight is that, if \mathcal{N} denotes the kernel of the singular operator, the subspaces should satisfy the *kernel-capturing property*

$$\mathcal{N} = \sum_i (V_i \cap \mathcal{N}). \quad (7.15)$$

For an $H(\operatorname{div})$ -based space V such as the MTW velocity space in the Stokes complex (3.1) or the AW stress space in the elasticity complex (4.11), one choice of relaxation method is given by the *vertex star iteration* [35], which is induced by the subspaces

$$V_i = \{v \in V \mid \operatorname{supp}(v) \subset K_i\}, \quad (7.16)$$

where K_i denotes the patch of cells in the mesh sharing vertex i . Given $v \in V$ with $\operatorname{div} v = 0$, where now $V \in \{V_h, \Sigma_h\}$ denotes one of the canonical MTW or AW spaces respectively, by exactness of the associated discrete complex when Ω is simply connected, we may write $\mathcal{C}\phi = v$ for some potential $\phi \in W$, where (W, \mathcal{C}) denote $(W_h/\mathbb{R}, \operatorname{curl})$ or $(Q_h, \operatorname{airy}_h)$ respectively. Let now $\{\phi_i\}_i$ denote a basis for W , and write $\phi = \sum_i c_i \phi_i$. Then a divergence-free decomposition of v is given by $v_i = c_i \mathcal{C}\phi_i$, since $v = \sum_i v_i$ and $v_i \in \mathcal{N}$ for each i , so it suffices to find subspaces V_i with $\mathcal{C}\phi_i \in V_i$ for each i . That the vertex star (7.16) fulfils this property follows from inspection of the basis functions of W .

In general, Schöberl's hypotheses also require that the splitting (7.14) be stable in the V -norm and that the kernel splitting (7.15) be stable in the energy norm induced by the Galerkin projection of the coercive form, which does not follow automatically from the exactness of the discrete complex; typically, such bounds hold for the infinite-dimensional spaces, and hold also for the discrete complex if bounded commuting projections exist.

8. Smoothers

To demonstrate the effectiveness of our mapping techniques in practical numerical simulations, it is necessary to consider preconditioners for the partial differential equations discretized by the MTW

and AW elements, and to exhibit composability of our implementation with the associated software stack [55]. As indicated in the previous subsection, the application of patch-based multigrid smoothers is natural for the $H(\text{div})$ -discretizing elements we consider in this paper.

Building on the work of Benzi and Olshanskii [15], Schöberl [80], and Hong et al. [44, 45] among others, Farrell and coauthors have successfully developed parameter- and mesh-robust preconditioners of augmented Lagrangian (AL) type, with specialized multigrid algorithms, for a host of nonlinear PDEs with saddle point structure [29, 30, 31, 33, 34, 35, 59, 87]. We illustrate the AL method for AW elements, the application to MTW being analogous. It consists of augmenting the Hellinger–Reissner Lagrangian (4.3) with a penalty term

$$\mathcal{H}_\alpha(\sigma, u) := \mathcal{H}(\sigma, u) + \frac{\alpha}{2} \int_\Omega \|\text{div}_h \sigma - f\|^2 dx \quad (8.1)$$

for $\alpha \geq 0$. The AL term penalizes deviation from the set constrained by the momentum balance (4.5b), but does not change the exact solution. Its more significant benefit is that it allows the control of the Schur complement of the discretized system, as we now explain. The stationarity condition of the augmented energy (8.1) gives rise to a saddle point system for the stress-displacement pair, which, by a block factorization, admits the well-known solution formula [14, Eq. 3.1]

$$\begin{bmatrix} A_\alpha & B^T \\ B & S \end{bmatrix}^{-1} = \begin{bmatrix} \mathbb{I} & -A_\alpha^{-1}B^T \\ & \mathbb{I} \end{bmatrix} \begin{bmatrix} A_\alpha^{-1} & \\ & S_\alpha^{-1} \end{bmatrix} \begin{bmatrix} \mathbb{I} & \\ & -BA_\alpha^{-1} \end{bmatrix} \mathbb{I}, \quad (8.2)$$

where A, B^T, B are discrete compliance, symmetric gradient, and divergence operators respectively, $S = -BA^{-1}B^T$ is the (in general, dense) Schur complement, and the subscript α denotes the same quantities but of the augmented system (so that $A_0 = A$ etc.). We wish to precondition this system for GMRES Krylov iterations. In analogy to the velocity-pressure Stokes problem, for which the Schur complement is spectrally equivalent to the viscosity-weighted pressure mass matrix [83], it can be shown (denoting by M_u the displacement mass matrix) that

$$\tilde{S}^{-1} := -\alpha M_u^{-1} \sim S_\alpha^{-1} \quad (8.3)$$

serves as an approximate inverse to the augmented Schur complement, at least for fixed mesh size, with the approximation improving as $\alpha \rightarrow \infty$. Preconditioning the Schur complement with (8.3) is however in tradeoff with the augmented stress solve A_α^{-1} ; the AL term in A_α has a large kernel consisting of the infinite-dimensional affine space of tensor fields with divergence f , rendering standard multigrid relaxation schemes ineffective. We propose the vertex star relaxation as an α -robust multigrid algorithm for this block.

Remark 8.1. All three of the finite element spaces we implement in this paper are non-nested under uniform refinement (in particular giving rise to non-nested bilinear forms); we employ the default prolongation operator of Firedrake, which involves (i) lossless projection of the coarse function onto a DG space of the same degree, (ii) lossless natural injection of the coarse projection, from the coarse grid to the fine DG space, (iii) projection from the fine DG space to the fine finite element space. We conjecture that multigrid convergence behaviour observed in Section 9 may be improved by the construction of specialized prolongation operators for each element which preserve, at least approximately, the kernel of the discrete divergence, in the style of [34, 35] (this being the other component of the parameter-robust multigrid framework of Schöberl), but we do not pursue this here.

Remark 8.2. Let $\tau_h \in \text{AW}^c(K)$ denote a local AW^c basis function dual to a vertex DOF for a given cell K . Because all other nodes vanish at τ_h , its full normal component vanishes along each edge, and its components have vanishing mean. Since $\varepsilon(\text{div } \tau_h)$ is a constant matrix,

$$\|\text{div } \tau_h\|_{0,K}^2 = \int_{\partial K} (\tau_h \mathbf{n}) \cdot (\text{div } \tau_h) ds - \int_K \tau_h : \varepsilon(\text{div } \tau_h) dx = 0; \quad (8.4)$$

thus, compared to AW^{nc} , the nodal AW^c basis functions arising from the “extra” vertex DOFs contribute only to the divergence-free subspace. It follows that similar multigrid transfer operators may work well for both these elements.

9. Numerical examples

For the case of affine rather than Piola transformations, we have described the inclusion of this theory in the Firedrake code stack in [56], and since this stack already understands Piola transformation the process is quite analogous. We must implement each new reference element in FIAT [53] and wrap it into FInAT [41]. The FInAT wrapper also requires a function to construct abstract syntax for the basis transformation in terms of callbacks provided by the form compiler to obtain symbols for geometric quantities such as Jacobians, and physical and reference normal and tangent vectors. The new elements must also be registered with UFL [1] and `tsfc` [42]. We now consider several test problems to validate our implementation and demonstrate its capabilities.

Linear systems in the manufactured solution examples, and on the coarsest grid in the multigrid examples, were solved by sparse LU factorization with MUMPS [2] and PETSc [13].

9.1. Mardal–Tai–Winther

We consider the following parameterized saddle point system: given $f, g \in L^2(\Omega; \mathbb{R}^2)$, $j \in H^{1/2}(\Gamma_D; \mathbb{R}^2)$, seek $(u, p) \in (H_{j,n,\Gamma_D}(\text{div}) \cap \epsilon H_{j,\Gamma_D}^1(\Omega; \mathbb{R}^2)) \times L^2(\Omega)$ such that

$$(I - \epsilon^2 \Delta) u + \nabla p = f \quad \text{in } \Omega, \quad (9.1a)$$

$$\text{div } u = g \quad \text{in } \Omega, \quad (9.1b)$$

$$u = j \quad \text{on } \Gamma_D, \quad (9.1c)$$

$$\epsilon^2 \nabla u \mathbf{n} - p \mathbf{n} = 0 \quad \text{on } \Gamma_N, \quad (9.1d)$$

in some domain $\Omega \subset \mathbb{R}^2$ with partitioned boundary $\Gamma_D \cup \Gamma_N$. Here, $\epsilon > 0$ is some positive parameter. Equations of this form arise in implicit timestepping for Stokes equations, for which ϵ corresponds to the square root of the timestep, and also in certain fractional-step methods for Navier–Stokes. For moderate ϵ , the equations are similar to the standard Stokes system and regular stable discretizations such as Taylor–Hood are readily applicable. However, as $\epsilon \searrow 0$, the equations degenerate to an $H(\text{div})$ Darcy system (equivalently, a mixed Poisson problem) and the stability and error estimates of standard H^1 methods degrade. The Mardal–Tai–Winther element, however, is stable for both $H(\text{div})$ and H^1 (though nonconforming) and for a scale of spaces in between; this is manifested in exhibiting stability and accuracy uniformly in ϵ for the system (9.1).

On the unit square $\Omega = (0, 1)^2$, with $\Gamma_N = \{y = 1\}$, $\Gamma_D = \partial\Omega \setminus \Gamma_N$, we approximate the manufactured solution $u(x, y) = (2^{1-y}, 0)^T$, $p(x, y) = \cos(\pi x) \cos(2\pi y)$. The essential boundary data (9.1c) is enforced weakly using the boundary DOFs and an interpolant of the given boundary data j . Denoting the resulting global constrained MTW space by $V_{h,j}$, and the global DG(0) pressure space by Q_h , we seek $(u_h, p_h) \in V_{h,j} \times Q_h$ such that

$$\begin{aligned} \int_{\Omega} u_h \cdot v_h + \epsilon^2 \nabla_h u_h : \nabla_h v_h - p_h \text{div } v_h \, dx &= \int_{\Omega} f \cdot v_h \, dx + \int_{\Gamma_N} K \mathbf{n} \cdot v_h \, ds \quad \forall v_h \in V_{h,0}, \\ - \int_{\Omega} q_h \text{div } u_h \, dx &= \int_{\Omega} g \cdot q_h \, dx \quad \forall q_h \in Q_h, \end{aligned} \quad (9.2)$$

where now f, g, j, K are defined by appropriate residuals of the manufactured solution. In order to show that the mapping techniques apply on general unstructured meshes, we perturb the interior vertices of the coarsest mesh as pictured in Figure 9.1; further refinements are obtained uniformly. In

PIOLA-MAPPED ELEMENTS

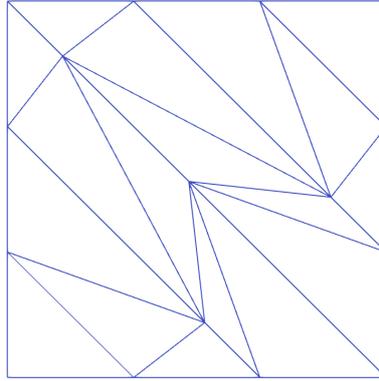


FIGURE 9.1. A warped mesh employed to check convergence under refinement.

Tables 9.1 and 9.2, we reproduce the approximately ϵ -independent experimental orders of convergence (EOCs) in both velocity and pressure proved in [64] for the MTW-DG(0) pair, for a representative range of parameters ϵ . Notice that the errors themselves are essentially ϵ -independent, and *a fortiori* so are the convergence rates, demonstrating the robustness of the element between these contrasting regimes.

TABLE 9.1. Errors and convergence rates in L^2 of the MTW velocity for the model problem (9.2). Here and below, N denotes the uniform refinement factor with respect to the original mesh.

$\epsilon \setminus N$	2^0	2^1	2^2	2^3	2^4	2^5	EOC
1	9.09×10^{-3}	2.45×10^{-3}	6.47×10^{-4}	1.68×10^{-4}	4.27×10^{-5}	1.08×10^{-5}	1.94
2^{-2}	8.79×10^{-3}	2.39×10^{-3}	6.29×10^{-4}	1.62×10^{-4}	4.13×10^{-5}	1.04×10^{-5}	1.94
2^{-4}	7.49×10^{-3}	2.17×10^{-3}	5.71×10^{-4}	1.46×10^{-4}	3.70×10^{-5}	9.32×10^{-6}	1.93
2^{-6}	6.81×10^{-3}	1.93×10^{-3}	5.25×10^{-4}	1.36×10^{-4}	3.46×10^{-5}	8.70×10^{-6}	1.92
2^{-8}	6.77×10^{-3}	1.88×10^{-3}	5.03×10^{-4}	1.31×10^{-4}	3.36×10^{-5}	8.51×10^{-6}	1.93
2^{-10}	6.77×10^{-3}	1.88×10^{-3}	5.02×10^{-4}	1.30×10^{-4}	3.30×10^{-5}	8.37×10^{-6}	1.93
0	6.77×10^{-3}	1.88×10^{-3}	5.02×10^{-4}	1.30×10^{-4}	3.30×10^{-5}	8.32×10^{-6}	1.93

TABLE 9.2. Errors and convergence rates in L^2 of pressure for (9.2).

$\epsilon \setminus N$	2^0	2^1	2^2	2^3	2^4	2^5	EOC
1	3.48×10^{-1}	1.83×10^{-1}	9.41×10^{-2}	4.73×10^{-2}	2.37×10^{-2}	1.18×10^{-2}	0.975
2^{-2}	3.08×10^{-1}	1.79×10^{-1}	9.35×10^{-2}	4.71×10^{-2}	2.36×10^{-2}	1.18×10^{-2}	0.941
2^{-4}	3.08×10^{-1}	1.79×10^{-1}	9.35×10^{-2}	4.71×10^{-2}	2.36×10^{-2}	1.18×10^{-2}	0.941
2^{-6}	3.08×10^{-1}	1.79×10^{-1}	9.35×10^{-2}	4.71×10^{-2}	2.36×10^{-2}	1.18×10^{-2}	0.941
2^{-8}	3.08×10^{-1}	1.79×10^{-1}	9.35×10^{-2}	4.71×10^{-2}	2.36×10^{-2}	1.18×10^{-2}	0.941
2^{-10}	3.08×10^{-1}	1.79×10^{-1}	9.35×10^{-2}	4.71×10^{-2}	2.36×10^{-2}	1.18×10^{-2}	0.941
0	3.08×10^{-1}	1.79×10^{-1}	9.35×10^{-2}	4.71×10^{-2}	2.36×10^{-2}	1.18×10^{-2}	0.941

We now demonstrate instead the robustness of our proposed multigrid solver for the MTW element applied to primal planar linear elasticity near the incompressible limit, a regime in which the MTW discretization will itself exhibit robust convergence; this example adapts the demonstration in [16]. We consider a simple cantilever beam $\Omega = [0, 25] \times [0, 1]$, fixed to a wall at the left-hand end Γ_D , stress-free on the rest of its boundary, in plane stress conditions, and subject only to the force of its own weight $f = (0, -10^{-3})^T$. Employing a Union-Jack-crossed mesh, and the associated MTW space

$V_{h,0}$, we seek $u_h \in V_{h,0}$ such that

$$\int_{\Omega} \mathbb{C}\varepsilon_h(u_h) : \varepsilon_h(v_h) \, dx = \int_{\Omega} f \cdot v_h \, dx \quad \forall v_h \in V_{h,0}, \tag{9.3}$$

where the shear modulus of the material is fixed at $\mu = 3.8 \times 10^4$ and the Poisson ratio is initialized at $\nu = 0.3$. In Table 9.3, for a discretization of 7.23×10^5 DOFs, we demonstrate robustness of the associated star relaxation, implemented via PCPATCH [32], as the Poisson ratio approaches the critical value $\frac{1}{2}$; the solution at highest Poisson ratio, enlarged for the sake of illustration, is pictured in Figure 9.2. This was carried out with relative ℓ^2 -norm tolerance 10^{-5} for the outermost Krylov iterations.

TABLE 9.3. Moderate and approximately constant Krylov iteration counts for the solution of (9.3) near the incompressible limit, after preconditioning via the vertex star relaxation described in subsection 7.2.

Poisson ratio	Krylov iterations
0.3	13
0.45	14
0.49	14
0.4999	14
0.4999999	14

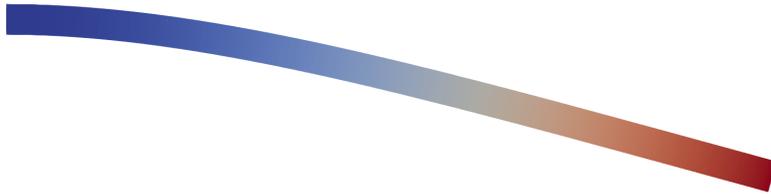


FIGURE 9.2. A nearly incompressible cantilever beam.

9.2. The Hellinger–Reissner principle

We now consider the canonical Hellinger–Reissner problem (6.1) and verify the convergence results proved by Arnold and Winther [11, Theorem 5.1][12, Theorem 4.1], summarized in Table 9.4. The approximation order m varies because in some cases, higher-order convergence may be obtained if the exact solution pair admits improved Sobolev regularity.

TABLE 9.4. Ranges of approximation order m by the AW elements of the elasticity variables in the L^2 -norm.

Variable	AW ^c	AW ^{nc}
σ	$1 \leq m \leq 3$	1
$\operatorname{div} \sigma$	$0 \leq m \leq 2$	$0 \leq m \leq 2$
u	$1 \leq m \leq 2$	1

On the unit square $\Omega = (0, 1)^2$, we consider the displacement from Bramwell et al. [18]: $u(x, y) = (\sin(\pi x) \sin(\pi y), \sin(\pi x) \sin(\pi y))^T$, which satisfies an unphysical homogeneous boundary condition,

and the stress field

$$\sigma(x, y) = \begin{bmatrix} \cos(\pi x) \cos(3\pi y) & y + 2 \cos(\frac{\pi}{2}x) \\ y + 2 \cos(\frac{\pi}{2}x) & -\sin(3\pi x) \cos(2\pi x) \end{bmatrix}. \quad (9.4)$$

In Tables 9.5 and 9.6, we check convergence of AW^{nc} in the L^2 -norms in the incompressible limit $\nu \nearrow \frac{1}{2}$, again on a perturbed mesh:

TABLE 9.5. Errors and convergence rates in the L^2 -norms using AW^{nc} for the model problem (6.1), with $\nu = 0.25, \mu = 1$.

N	u error	u EOC	σ error	σ EOC	$\text{div}_h(\sigma)$ error	$\text{div}_h(\sigma)$ EOC
2^0	1.05×10^{-1}	–	5.51×10^{-1}	–	8.32×10^{-1}	–
2^1	2.78×10^{-2}	1.91	2.54×10^{-2}	1.12	2.68×10^{-1}	1.64
2^2	7.07×10^{-3}	1.98	1.18×10^{-1}	1.10	7.15×10^{-2}	1.91
2^3	1.77×10^{-3}	2.00	5.76×10^{-2}	1.04	1.82×10^{-2}	1.98
2^4	4.43×10^{-4}	2.00	2.85×10^{-2}	1.02	4.56×10^{-3}	1.99
2^5	1.11×10^{-4}	2.00	1.42×10^{-2}	1.00	1.14×10^{-3}	2.00

TABLE 9.6. Errors and convergence rates with AW^{nc} near the incompressible limit $\nu = 0.4999999, \mu = 1$.

N	u error	u EOC	σ error	σ EOC	$\text{div}_h(\sigma)$ error	$\text{div}_h(\sigma)$ EOC
2^0	1.05×10^{-1}	–	6.53×10^{-1}	–	8.32×10^{-1}	–
2^1	2.78×10^{-2}	1.91	3.07×10^{-1}	1.09	2.68×10^{-1}	1.64
2^2	7.05×10^{-3}	1.98	1.49×10^{-1}	1.04	7.15×10^{-2}	1.91
2^3	1.77×10^{-3}	2.00	7.38×10^{-2}	1.02	1.82×10^{-2}	1.98
2^4	4.42×10^{-4}	2.00	3.66×10^{-2}	1.01	4.56×10^{-3}	1.99
2^5	1.11×10^{-4}	2.00	1.83×10^{-2}	1.00	1.14×10^{-3}	2.00

Note that the observed order of convergence in the displacement is one higher than proved, but 2^{nd} -order convergence for the displacement is proved for the 3D analogue of AW^{nc} , the nonconforming Arnold–Awanou–Winther element [5], by applying a duality argument in the case of full elliptic regularity $(\sigma, u) \in H^1(\Omega; \mathbb{S}) \times H^2(\Omega; \mathbb{R}^d)$.

Tables 9.7–9.8 exhibit convergence behaviour of the AW^{c} element, applied to $u(x, y) = (-e^{\sin(\frac{\pi}{2}y)}, 3 \cos(\pi x))^T$, and exact stress $\sigma = \mathbb{C}\varepsilon(u)$.

TABLE 9.7. Errors and convergence rates using AW^{c} with $\nu = 0.25, \mu = 1$.

N	u error	u EOC	σ error	σ EOC	$\text{div}(\sigma)$ error	$\text{div}(\sigma)$ EOC
2^0	1.50×10^{-1}	–	1.16×10^1	–	1.47×10^0	–
2^1	3.81×10^{-2}	1.97	1.60×10^{-2}	2.85	3.75×10^{-1}	1.97
2^2	9.57×10^{-3}	1.99	2.09×10^{-3}	2.94	9.44×10^{-2}	1.99
2^3	2.40×10^{-3}	2.00	2.65×10^{-4}	2.98	2.36×10^{-2}	2.00
2^4	5.99×10^{-4}	2.00	3.34×10^{-5}	2.99	5.91×10^{-3}	2.00
2^5	1.50×10^{-4}	2.00	4.19×10^{-6}	3.00	1.48×10^{-3}	2.00

TABLE 9.8. AW^{c} near the incompressible limit $\nu = 0.4999999, \mu = 1$.

N	u error	u EOC	σ error	σ EOC	$\text{div}(\sigma)$ error	$\text{div}(\sigma)$ EOC
2^0	1.49×10^{-1}	–	1.70×10^{-1}	–	1.47×10^0	–
2^1	3.81×10^{-2}	1.97	1.99×10^{-2}	3.09	3.75×10^{-1}	1.97
2^2	9.57×10^{-3}	1.99	2.31×10^{-3}	3.12	9.44×10^{-2}	1.99
2^3	2.40×10^{-3}	2.00	2.80×10^{-4}	3.05	2.36×10^{-2}	2.00
2^4	5.99×10^{-4}	2.00	3.46×10^{-5}	3.01	5.91×10^{-3}	2.00
2^5	1.50×10^{-4}	2.00	4.32×10^{-6}	3.00	1.48×10^{-3}	2.00

We manifestly see robustness in the incompressible regime for both elements. The comparative magnitudes of the errors show that the AW^c element provides a markedly more accurate approximation to the stress field. Note also that convergence behaviour in the divergence of the stress is the same for both elements, since their divergences have the same degree; this is also consistent with the “extra” AW^c vertex basis functions being solenoidal as observed in Remark 8.2.

Having validated both AW elements, we consider a more complex example from the PhD thesis of Li [62] which moreover includes traction conditions. The domain, pictured in Figure 9.3, consists of the rectangle $\Omega = [0, 3] \times [0, 1]$ with three disks removed, occupied by a material assumed to be isotropic and homogeneous with $\nu = 0.2$ and Young’s modulus $E = 10$:

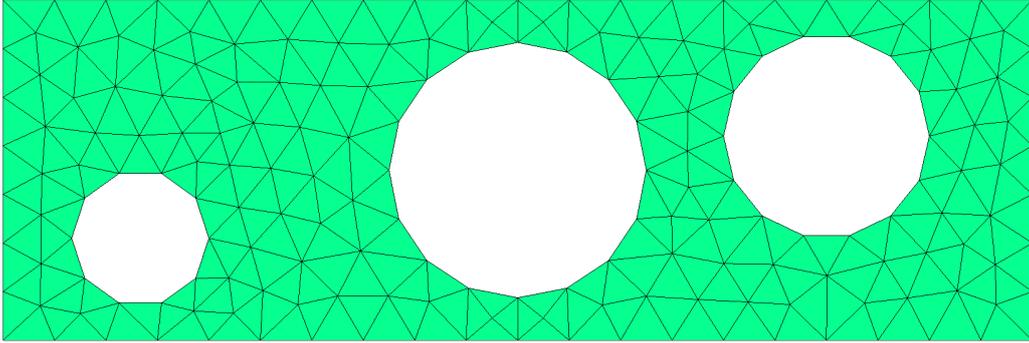


FIGURE 9.3. Domain with a coarser initial mesh than that employed in [62] so that multigrid may be performed.

The prescribed displacement is fixed $(0, 0)^T$ at the left-hand end and compressed $(-1, 0)^T$ at the right end, which together form Γ_D , with a stress-free condition $\sigma \mathbf{n} = 0$ on $\Gamma_N = \partial\Omega \setminus \Gamma_D$ given by the top, bottom, and the boundaries of the holes; there is no external force f . We combine the Nitsche and augmented Lagrangian penalties, seeking critical points (σ_h, u_h) of

$$\mathcal{H}_{h,\gamma,\alpha}(\sigma_h, u_h) := \mathcal{H}_{h,\gamma}(\sigma_h, u_h) + \frac{\alpha}{2} \int_{\Omega} \|\operatorname{div}_h \sigma_h\|^2 dx. \quad (9.5)$$

Due to the Nitsche boundary terms, which at present cannot be treated with PCPATCH, the application of the vertex star iteration necessitated the use of PCASM. This was carried out with residual ℓ^2 -norm tolerance 10^{-9} for the outermost Krylov iterations. Table 9.9 exhibits the behaviour of GMRES for AW^c , with multigrid for the augmented stress block employing vertex star relaxation with Chebyshev smoother on each multigrid level, with fixed Nitsche and AL parameters $\gamma = 100, \alpha = 1$. We verify in Figure 9.4, which is colored by the size of the shear stress $\|\operatorname{dev} \sigma_h\|$, that the AW^c solution on the finest mesh, with almost five million degrees of freedom, is free from numerical artifacts. Figure 9.5 shows convergence of the traction residual to zero in $L^2(\Gamma_N; \mathbb{R}^2)$ as $h \rightarrow 0$ using AW^c with $\alpha = 1$, and AW^{nc} (using LU factorization) with $\alpha = 0$ respectively, for various values of the Nitsche parameter γ . Although we have proved convergence of the Nitsche penalty for any $\gamma \geq 1$, in practice we find the solver most effective at $\gamma = \mathcal{O}(100)$, and that increasing γ further will enforce the traction condition more strongly (if desired), at the cost of a more ill-conditioned system.

Remark 9.1. The observed order of convergence for the traction residual in Figure 9.5 is higher than that predicted by Proposition 6.6, which itself makes an artificial regularity assumption on the stress field, and which we therefore conjecture could be improved, for example via duality arguments.

PIOLA-MAPPED ELEMENTS

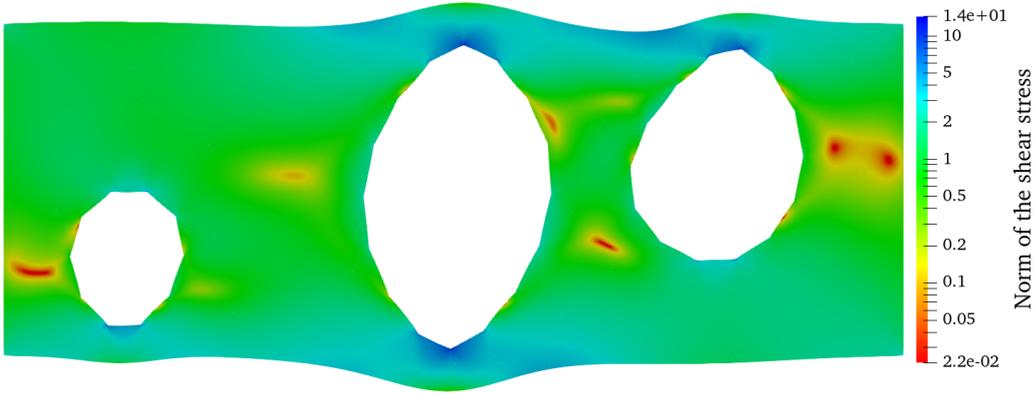
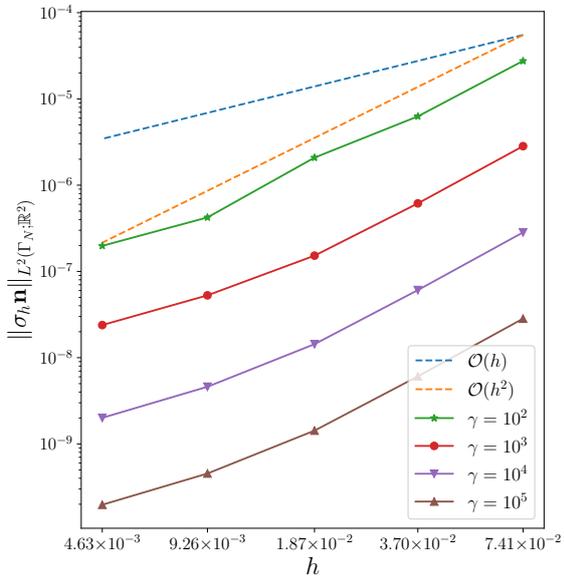


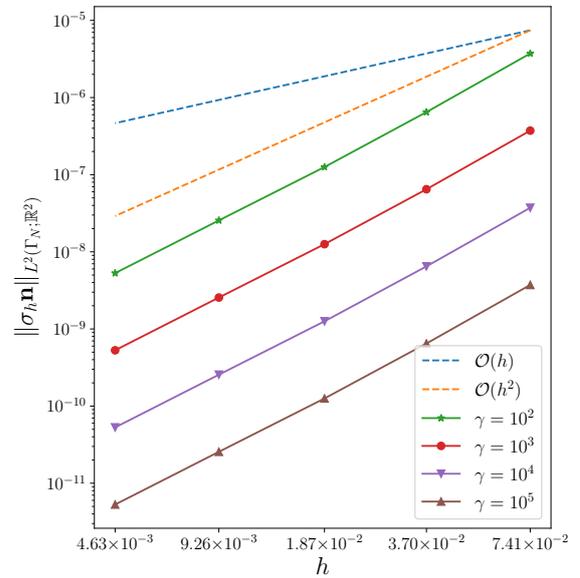
FIGURE 9.4. A traction-free condition except at both ends.

TABLE 9.9. Outer Krylov iteration counts using the vertex star relaxation for the AW^c element.

N	DOFs	Krylov iterations
2^1	2.01×10^4	12
2^2	7.90×10^4	12
2^3	3.13×10^5	12
2^4	1.25×10^6	12
2^5	4.98×10^6	11



AW^c with $\alpha = 1$



AW^{nc} with $\alpha = 0$

FIGURE 9.5. Convergence to the traction-free condition in $L^2(\Gamma_N; \mathbb{R}^2)$ for the conforming and nonconforming AW elements, with and without AL penalization respectively.

9.3. Code availability

The Mardal–Tai–Winther and conforming and nonconforming Arnold–Winther elements, with the modifications suggested by subsection 5.4, were incorporated into the main branch of Firedrake as part of this work. For reproducibility, the exact software versions used to generate the numerical results in this paper are archived in <https://zenodo.org/record/5596313> [90]; the code, and scripts for the associated plots, are available at <https://bitbucket.org/FAznaran/piola-mapped>.

10. Conclusion

We have generalized Piola transformation theory to incorporate non-standard $H(\text{div})$ elements, developing techniques for representative finite elements discretizing $H(\text{div})$ and $H(\text{div}; \mathbb{S})$ in two dimensions: the Mardal–Tai–Winther element for Stokes–Darcy flow, and two exotic, symmetry-enforcing elements for elastic stress due to Arnold and Winther. Numerical experiments verify the accuracy of implementations which are newly enabled by our approach. All elements have been implemented in the publicly available Firedrake library, within which we have demonstrated the composability of our implementations with existing patch-based smoothers. Of independent interest is the application of Nitsche’s method to dual mixed problems; further investigation is merited by the interaction between the Nitsche and augmented Lagrangian penalties, whose efficacy when combined together we have observed numerically when applied to Arnold–Winther elements for linear elasticity.

We emphasize that our theory aims to demonstrate that unusual or non-standard elements, with desirable features but perhaps complicated transformation properties, can be used in an inexpensive, composable, and automated way, rather than to advocate for the use of MTW, AW^c , or AW^{nc} specifically. In particular, the 3D analogue of the AW spaces, namely the conforming and nonconforming Arnold–Awanou–Winther elements [4, 5], are of local dimension 162 and 42 respectively, and are presented in the understanding that “[t]he complexity of the elements may very well limit their practical significance” [4, p. 1231]. A cheaper alternative is provided by the conforming Hu–Zhang element, which in 3D is of dimension only 48, while many nonconforming efforts in 3D are rectangular [47, 89]; in keeping with Remark 4.1, we conjecture that the mapping properties of these 3D elements would be analogous to our analysis in Section 5.

Acknowledgements

The authors are grateful to Endre Süli for comments on Section 6, to Kaibo Hu for comments on subsection 7.1, and to the two anonymous referees for their valuable suggestions.

References

- [1] M. S. Alnæs, A. Logg, K. B. Ølgaard, M. E. Rognes, and G. N. Wells. Unified Form Language: a Domain-Specific Language for weak formulations of partial differential equations. *ACM Trans. Math. Softw.*, 40(2):1–37, 2014.
- [2] P. R. Amestoy, I. S. Duff, J.-Y. L’Excellent, and J. Koster. A fully asynchronous multifrontal solver using distributed dynamic scheduling. *SIAM J. Matrix Anal. Appl.*, 23(1):15–41, 2001.
- [3] D. N. Arnold. *Finite element exterior calculus*. Number 93 in CBMS-NSF Regional Conference Series in Applied Mathematics. Society for Industrial and Applied Mathematics, 2018.
- [4] D. N. Arnold, G. Awanou, and R. Winther. Finite elements for symmetric tensors in three dimensions. *Math. Comput.*, 77(263):1229–1251, 2008.

- [5] D. N. Arnold, G. Awanou, and R. Winther. Nonconforming tetrahedral mixed finite elements for elasticity. *Math. Models Methods Appl. Sci.*, 24(04):783–796, 2014.
- [6] D. N. Arnold, F. Brezzi, and J. Douglas. PEERS: A new mixed finite element for plane elasticity. *Japan J. Appl. Math.*, 1(2):347–367, 1984.
- [7] D. N. Arnold, R. S. Falk, and R. Winther. Multigrid in $H(\text{div})$ and $H(\text{curl})$. *Numer. Math.*, 85(2):197–217, 2000.
- [8] D. N. Arnold, R. S. Falk, and R. Winther. Finite element exterior calculus, homological techniques, and applications. *Acta Numer.*, 15:1–155, 2006.
- [9] D. N. Arnold, R. S. Falk, and R. Winther. Mixed finite element methods for linear elasticity with weakly imposed symmetry. *Math. Comput.*, 76(260):1699–1723, 2007.
- [10] D. N. Arnold, R. S. Falk, and R. Winther. Geometric decompositions and local bases for spaces of finite element differential forms. *Comput. Methods Appl. Mech. Eng.*, 198(21-26):1660–1672, 2009.
- [11] D. N. Arnold and R. Winther. Mixed finite elements for elasticity. *Numer. Math.*, 92(3):401–419, 2002.
- [12] D. N. Arnold and R. Winther. Nonconforming mixed elements for elasticity. *Math. Models Methods Appl. Sci.*, 13(03):295–307, 2003.
- [13] S. Balay, S. Abhyankar, M. F. Adams, J. Brown, P. Brune, K. Buschelman, L. Dalcin, V. Eijkhout, W. D. Gropp, D. Karpeyev, D. Kaushik, M. G. Knepley, D. May, L. C. McInnes, R. T. Mills, T. Munson, K. Rupp, P. Sanan, B. F. Smith, S. Zampini, H. Zhang, and H. Zhang. PETSc Users Manual. Technical report, Argonne National Laboratory, 2019. ANL-95/11 - Revision 3.11.
- [14] M. Benzi, G. H. Golub, and J. Liesen. Numerical solution of saddle point problems. *Acta Numer.*, 14:1–137, 2005.
- [15] M. Benzi and M. A. Olshanskii. An augmented Lagrangian-based approach to the Oseen problem. *SIAM J. Sci. Comput.*, 28(6):2095–2113, 2006.
- [16] J. Bleyer. *Numerical Tours of Computational Mechanics with FEniCS*, 2018. <https://comet-fenics.readthedocs.io>.
- [17] D. Boffi, F. Brezzi, and M. Fortin. *Mixed finite element methods and applications*, volume 44 of *Springer Series in Computational Mathematics*. Springer, 2013.
- [18] J. Bramwell, L. Demkowicz, J. Gopalakrishnan, and W. Qiu. A locking-free hp DPG method for linear elasticity with symmetric stresses. *Numer. Math.*, 122(4):671–707, 2012.
- [19] F. Brezzi, J. Douglas, M. Fortin, and L. D. Marini. Efficient rectangular mixed finite elements in two and three space variables. *ESAIM, Math. Model. Numer. Anal.*, 21(4):581–604, 1987.
- [20] F. Brezzi, J. Douglas, and L. D. Marini. Two families of mixed finite elements for second order elliptic problems. *Numer. Math.*, 47(2):217–235, 1985.
- [21] E. Burman and R. Puppi. Two mixed finite element formulations for the weak imposition of the Neumann boundary conditions for the Darcy flow. *J. Numer. Math.*, 30(2):141–162, 2022.
- [22] Z. Cai and G. Starke. Least-squares methods for linear elasticity. *SIAM J. Numer. Anal.*, 42(2):826–842, 2004.
- [23] C. Carstensen, M. Eigel, and J. Gedicke. Computational competition of symmetric mixed FEM in linear elasticity. *Comput. Methods Appl. Mech. Eng.*, 200(41-44):2903–2915, 2011.
- [24] C. Carstensen, D. Gallistl, and M. Schedensack. L^2 best approximation of the elastic stress in the Arnold–Winther FEM. *IMA J. Numer. Anal.*, 36(3):1096–1119, 2016.
- [25] C. Carstensen, D. Günther, J. Reininghaus, and J. Thiele. The Arnold–Winther mixed FEM in linear elasticity. Part I: Implementation and numerical verification. *Comput. Methods Appl. Mech. Eng.*, 197(33-40):3014–3023, 2008.

- [26] L. Chen and X. Huang. Geometric decompositions of div-conforming finite element tensors. <https://arxiv.org/abs/2112.14351>, 2021.
- [27] S. H. Christiansen, J. Hu, and K. Hu. Nodal finite element de Rham complexes. *Numer. Math.*, 139(2):411–446, 2018.
- [28] G. Danisch and G. Starke. First-order system least-squares for Darcy–Stokes flow. *SIAM J. Numer. Anal.*, 45(2):731–745, 2007.
- [29] P. E. Farrell, L. F. Gatica, B. P. Lamichhane, R. Oyarzúa, and R. Ruiz-Baier. Mixed Kirchhoff stress–displacement–pressure formulations for incompressible hyperelasticity. *Comput. Methods Appl. Mech. Eng.*, 374:113562, 2021.
- [30] P. E. Farrell and P. A. Gazca-Orozco. An augmented Lagrangian preconditioner for implicitly-constituted non-Newtonian incompressible flow. *SIAM J. Sci. Comput.*, 42(6):B1329–B1349, 2020.
- [31] P. E. Farrell, P. A. Gazca-Orozco, and E. Süli. Finite element approximation and preconditioning for anisothermal flow of implicitly-constituted non-Newtonian flow. *Math. Comput.*, 91(334):659–697, 2021.
- [32] P. E. Farrell, M. G. Knepley, F. Wechsung, and L. Mitchell. PCPATCH: software for the topological construction of multigrid relaxation methods. *ACM Trans. Math. Softw.*, 47(3):1–22, 2021.
- [33] P. E. Farrell, L. Mitchell, L. R. Scott, and F. Wechsung. A Reynolds-robust preconditioner for the Scott–Vogelius discretization of the stationary incompressible Navier–Stokes equations. *SMAI J. Comput. Math.*, 7:75–96, 2021.
- [34] P. E. Farrell, L. Mitchell, L. R. Scott, and F. Wechsung. Robust multigrid methods for nearly incompressible elasticity using macro elements. *IMA J. Numer. Anal.*, 2022.
- [35] P. E. Farrell, L. Mitchell, and F. Wechsung. An augmented Lagrangian preconditioner for the 3D stationary incompressible Navier–Stokes equations at high Reynolds number. *SIAM J. Sci. Comput.*, 41(5):A3073–A3096, 2019.
- [36] T. Frankel. *The geometry of physics: An introduction*. Cambridge University Press, 2nd edition, 2003.
- [37] J. Gedicke and A. Khan. Arnold–Winther mixed finite elements for Stokes eigenvalue problems. *SIAM J. Sci. Comput.*, 40(5):A3449–A3469, 2018.
- [38] J. S. Hale, M. Brunetti, S. P. A. Bordas, and C. Maurini. Simple and extensible plate and shell finite element models through automatic code generation tools. *Computers & Structures*, 209:163–181, 2018.
- [39] K. Hellan. Analysis of elastic plates in flexure by a simplified finite element method. *Acta polytechn. Scand., Civil Engin. Building Construct. Ser.*, 46:1–29, 1967. Finnish Academy of Technology, Tekniikentie 12, FIN-02150 Espoo, Finland.
- [40] L. R. Herrmann. Finite-element bending analysis for plates. *J. Eng. Mech. Div.*, 93(5):13–26, 1967.
- [41] M. Homolya, R. C. Kirby, and D. A. Ham. Exposing and exploiting structure: Optimal code generation for high-order finite element methods. <https://arxiv.org/abs/1711.02473>, 2017.
- [42] M. Homolya, L. Mitchell, F. Luporini, and D. A. Ham. TSFC: A structure-preserving form compiler. *SIAM J. Sci. Comput.*, 40(3):C401–C428, 2018.
- [43] Q. Hong, J. Hu, L. Ma, and J. Xu. An extended Galerkin analysis for linear elasticity with strongly symmetric stress tensor. <https://arxiv.org/abs/2002.11664>, 2020.
- [44] Q. Hong and J. Kraus. Uniformly stable discontinuous Galerkin discretization and robust iterative solution methods for the Brinkman problem. *SIAM J. Numer. Anal.*, 54(5):2750–2774, 2016.
- [45] Q. Hong, J. Kraus, J. Xu, and L. Zikatanov. A robust multigrid method for discontinuous Galerkin discretizations of Stokes and linear elasticity equations. *Numer. Math.*, 132:23–49, 2015.
- [46] J. Hu. Finite element approximations of symmetric tensors on simplicial grids in \mathbb{R}^n : The higher order case. *J. Comput. Math.*, 33(3):283–296, 2015.

- [47] J. Hu. A new family of efficient conforming mixed finite elements on both rectangular and cuboid meshes for linear elasticity in the symmetric formulation. *SIAM J. Numer. Anal.*, 53(3):1438–1463, 2015.
- [48] J. Hu and S. Zhang. A family of conforming mixed finite elements for linear elasticity on triangular grids. <https://arxiv.org/abs/1406.7457>, 2015.
- [49] J. Hu and S. Zhang. Finite element approximations of symmetric tensors on simplicial grids in \mathbb{R}^n : The lower order case. *Math. Models Methods Appl. Sci.*, 26(09):1649–1669, 2016.
- [50] V. John, A. Linke, C. Merdon, M. Neilan, and L. G. Rebholz. On the divergence constraint in mixed finite element methods for incompressible flows. *SIAM Rev.*, 59(3):492–544, 2017.
- [51] C. Johnson. On the convergence of a mixed finite-element method for plate bending problems. *Numer. Math.*, 21(1):43–62, 1973.
- [52] T. Karper, K.-A. Mardal, and R. Winther. Unified finite element discretizations of coupled Darcy–Stokes flow. *Numerical Methods for Partial Differential Equations: An International Journal*, 25(2):311–326, 2009.
- [53] R. C. Kirby. Algorithm 839: FIAT, A New Paradigm for Computing Finite Element Basis Functions. *ACM Trans. Math. Softw.*, 30(4):502–516, 2004.
- [54] R. C. Kirby. A general approach to transforming finite elements. *SMAI J. Comput. Math.*, 4:197–224, 2018.
- [55] R. C. Kirby and L. Mitchell. Solver composition across the PDE/linear algebra barrier. *SIAM J. Sci. Comput.*, 40(1):C76–C98, 2018.
- [56] R. C. Kirby and L. Mitchell. Code generation for generally mapped finite elements. *ACM Trans. Math. Softw.*, 45(4):1–23, 2019.
- [57] J. Könnö, D. Schötzau, and R. Stenberg. Mixed finite element methods for problems with Robin boundary conditions. *SIAM J. Numer. Anal.*, 49(1):285–308, 2011.
- [58] M. Kuchta, K.-A. Mardal, and M. Mortensen. On the singular Neumann problem in linear elasticity. *Numer. Linear Algebra Appl.*, 26(1), 2019.
- [59] F. Laakmann, P. E. Farrell, and L. Mitchell. An augmented Lagrangian preconditioner for the magneto-hydrodynamics equations at high Reynolds and coupling numbers. *SIAM J. Sci. Comput.*, 44(4):B1018–B1044, 2022.
- [60] H. P. Langtangen and A. Tveito. *Advanced topics in computational partial differential equations: Numerical methods and Diffpack programming*, volume 33 of *Lecture Notes in Computational Science and Engineering*. Springer, 2012.
- [61] J. J. Lee. Robust three-field finite element methods for Biot’s consolidation model in poroelasticity. *BIT Numer. Math.*, 58(2):347–372, 2018.
- [62] L. Li. *Regge finite elements with applications in solid mechanics and relativity*. PhD thesis, University of Minnesota, 2018. <http://www-users.math.umn.edu/~arnold/papers/LiThesis.pdf>.
- [63] A. Logg, K.-A. Mardal, and G. N. Wells, editors. *Automated solution of differential equations by the finite element method: The FEniCS book*, volume 84 of *Lecture Notes in Computational Science and Engineering*. Springer, 2012.
- [64] K.-A. Mardal, X.-C. Tai, and R. Winther. A robust finite element method for Darcy–Stokes flow. *SIAM J. Numer. Anal.*, 40(5):1605–1631, 2002.
- [65] K.-A. Mardal and R. Winther. An observation on Korn’s inequality for nonconforming finite element methods. *Math. Comput.*, 75(253):1–6, 2006.
- [66] L. A. Mihai and A. Goriely. Numerical simulation of shear and the Poynting effects by the finite element method: An application of the generalised empirical inequalities in non-linear elasticity. *Int. J. Non-Linear Mech.*, 49:1–14, 2013.
- [67] P. Monk. *Finite Element Methods for Maxwell’s Equations*. Numerical Mathematics and Scientific Computation. Oxford University Press, 2003.

- [68] B. Müller, G. Starke, A. Schwarz, and J. Schröder. A first-order system least squares method for hyperelasticity. *SIAM J. Sci. Comput.*, 36(5):B795–B816, 2014.
- [69] S. Nicaise, K. Witowski, and B. Wohlmuth. An a posteriori error estimator for the Lamé equation based on $H(\text{div})$ -conforming stress approximations. *IMA J. Numer. Anal.*, 28:331–353, 2008.
- [70] T. Nilssen, X.-C. Tai, and R. Winther. A robust nonconforming H^2 -element. *Math. Comput.*, 70(234):489–505, 2001.
- [71] J. Nitsche. Über ein Variationsprinzip zur Lösung von Dirichlet-Problemen bei Verwendung von Teilräumen, die keinen Randbedingungen unterworfen sind. In *Abhandlungen aus dem mathematischen Seminar der Universität Hamburg*, volume 36, pages 9–15. Springer, 1971.
- [72] S. Olver and Y. Xu. Orthogonal structure on a wedge and on the boundary of a square. *Found. Comput. Math.*, 19(3):561–589, 2019.
- [73] J. E. Pasciak and Y. Wang. A multigrid preconditioner for the mixed formulation of linear plane elasticity. *SIAM J. Numer. Anal.*, 44(2):478–493, 2006.
- [74] D. Pauly and W. Zulehner. The divDiv-complex and applications to biharmonic equations. *Appl. Anal.*, 99(9):1579–1630, 2020.
- [75] R. Rannacher and S. Turek. Simple nonconforming quadrilateral Stokes element. *Numer. Methods Partial Differ. Equations*, 8(2):97–111, 1992.
- [76] F. Rathgeber, D. A. Ham, L. Mitchell, M. Lange, F. Luporini, A. T. T. McRae, G.-T. Bercea, G. R. Markall, and P. H. J. Kelly. Firedrake: Automating the finite element method by composing abstractions. *ACM Trans. Math. Softw.*, 43(3):1–27, 2016.
- [77] P.-A. Raviart and J.-M. Thomas. A mixed finite element method for 2nd order elliptic problems. In *Mathematical Aspects of Finite Element Methods*, pages 292–315. Springer, 1977.
- [78] R. Riedlbeck, D. A. Di Pietro, A. Ern, S. Granet, and K. Kazymyrenko. Stress and flux reconstruction in Biot’s poro-elasticity problem with application to a posteriori error analysis. *Comput. Math. Appl.*, 73(7):1593–1610, 2017.
- [79] M. E. Rognes, R. C. Kirby, and A. Logg. Efficient assembly of $H(\text{div})$ and $H(\text{curl})$ conforming finite elements. *SIAM J. Sci. Comput.*, 31(6):4130–4151, 2009.
- [80] J. Schöberl. *Robust multigrid methods for parameter dependent problems*. PhD thesis, Johannes Kepler University, 1999. <https://www.asc.tuwien.ac.at/~schoeberl/wiki/publications/diss.pdf>.
- [81] J. Schöberl and A. S. Sinwel. Tangential-Displacement and Normal-Normal-Stress continuous mixed finite Elements for elasticity. Technical report, Radon Institute for Computational and Applied Mathematics, 2007. <https://ricamwww.ricam.oeaw.ac.at/files/reports/07/rep07-10.pdf>.
- [82] D. Shapero. Nitsche’s method. <https://shapero.xyz/posts/nitsches-method/>, 2019. Accessed 2020.11.16.
- [83] D. Silvester and A. Wathen. Fast iterative solution of stabilised Stokes systems Part II: Using general block preconditioners. *SIAM J. Numer. Anal.*, 31(5):1352–1367, 1994.
- [84] X.-C. Tai and R. Winther. A discrete de Rham complex with enhanced smoothness. *Calcolo*, 43(4):287–306, 2006.
- [85] Y. Wang. *Preconditioning for the mixed formulation of linear plane elasticity*. PhD thesis, Texas A&M University, 2004. <https://oaktrust.library.tamu.edu/bitstream/handle/1969.1/2781/etd-tamu-2004B-MATH-Wang.pdf>.
- [86] S. Wu and J. Xu. Nonconforming finite element spaces for $2m$ th order partial differential equations on \mathbb{R}^n simplicial grids when $m = n + 1$. *Math. Comput.*, 88(316):531–551, 2019.
- [87] J. Xia, P. E. Farrell, and F. Wechsung. Augmented Lagrangian preconditioners for the Oseen–Frank model of nematic and cholesteric liquid crystals. *BIT Numer. Math.*, pages 1–38, 2021.

- [88] J. Xu. Iterative methods by space decomposition and subspace correction. *SIAM Rev.*, 34(4):581–613, 1992.
- [89] S-Y Yi. Nonconforming mixed finite element methods for linear elasticity using rectangular elements in two and three dimensions. *Calcolo*, 42(2):115–133, 2005.
- [90] Software used in ‘Transformations for Piola-mapped elements’, 2021. <https://doi.org/10.5281/zenodo.5596313>.