SANDIA REPORT
SAND2020-9980
Printed September, 2020

Sandia National Laboratories

Regression Based Approach for Robust Finite Element Analysis on Arbitrary Grids: LDRD Final Report

Paul Kuberry, Pavel Bochev, Jacob Koester, Nathaniel Trask

Prepared by
Sandia National Laboratories
Albuquerque, New Mexico 87185
Livermore, California 94550

Issued by Sandia National Laboratories, operated for the United States Department of Energy by National Technology & Engineering Solutions of Sandia, LLC.

Available to DOE and DOE contractors from

U.S. Department of Energy
Office of Scientific and Technical Information
P.O. Box 62
Oak Ridge, TN 37831

Telephone: (865) 576-8401
Facsimile: (865) 576-5728
E-Mail: reports@osti.gov
Online ordering: http://www.osti.gov/scitech

Available to the public from

U.S. Department of Commerce
National Technical Information Service
5301 Shawnee Road
Alexandria, VA 22312

Telephone: (800) 553-6847
Facsimile: (703) 605-6900
E-Mail: orders@ntis.gov
Online order: https://classic.ntis.gov/help/order-methods

ABSTRACT

This report summarizes the work performed under a one-year LDRD project aiming to enable accurate and robust numerical simulation of partial differential equations for meshes that are of poor quality.

Traditional finite element methods use the mesh to both discretize the geometric domain and to define the finite element shape functions. The latter creates a dependence between the quality of the mesh and the properties of the finite element basis that may adversely affect the accuracy of the discretized problem. In this project, we propose a new approach for defining finite element shape functions that breaks this dependence and separates mesh quality from the discretization quality. At the core of the approach is a meshless definition of the shape functions, which limits the purpose of the mesh to representing the geometric domain and integrating the basis functions without having any role in their approximation quality.

The resulting non-conforming space can be utilized within a standard discontinuous Galerkin framework providing a rigorous foundation for solving partial differential equations on low-quality meshes. We present a collection of numerical experiments demonstrating our approach in a wide range of settings: strongly coercive elliptic problems, linear elasticity in the compressible regime, and the stationary Stokes problem.

We demonstrate convergence for all problems and stability for element pairs for problems which usually require inf-sup compatibility for conforming methods, also referring to a minor modification possible through the symmetric interior penalty Galerkin framework for stabilizing element pairs that would otherwise be traditionally unstable. Mesh robustness is particularly critical for elasticity, and we provide an example that our approach provides a greater than 5x improvement in accuracy and allows for taking an 8x larger stable timestep for a highly deformed mesh, compared to the continuous Galerkin finite element method.

The report concludes with a brief summary of ongoing projects and collaborations that utilize or extend the products of this work.

ACKNOWLEDGMENT

Supported by the Laboratory Directed Research and Development program at Sandia National Laboratories, a multimission laboratory managed and operated by National Technology and Engineering Solutions of Sandia LLC, a wholly owned subsidiary of Honeywell International Inc. for the U.S. Department of Energy's National Nuclear Security Administration under contract DE-NA0003525.

CONTENTS

LIST OF FIGURES

LIST OF TABLES

NOMENCLATURE

Table 0-1.

Abbreviation	Definition
REBAR	regression based approach for robust finite element analysis on arbitrary grids
PDE	partial differential equation
FE	finite element
FEA	finite element analysis
FEM	finite element method
GMLS	generalized moving least squares
CG	continuous Galerkin
DG	discontinuous Galerkin
IP	interior penalty
SIPG	symmetric interior penalty Galerkin
VMS	variational multiscale

1. INTRODUCTION

The classical definition of a finite element (FE) space [10] as a triple comprising (i) a geometric region (the element), (ii) a unisolvent set of functionals (the degrees of freedom), and (iii) a polynomial space, ties together the shape of the element and the properties of the resulting basis.

Thanks to the pioneering work of Babuska [4] it is now well understood that poorly shaped elements may affect adversely the quality of the resulting FE discretization. In particular, [4] established that the minimum angle condition, long thought to be of critical importance, is in fact not essential and that one should instead avoid elements that have angles close to π.

This interdependence between element shape and basis quality propagates through the entire foundation of finite element analysis (FEA) by affecting key properties such as inverse inequalities, condition numbers, and convergence estimates. As a result, the general rule of the thumb so far has been to avoid as much as possible bad elements such as high-aspect or "sliver" elements [20].

Unfortunately, this is not always possible in many engineering and scientific applications that require the meshing of complex geometric domains because automatic generation of high-quality grids remains a challenge. Currently, hexahedral grids can deliver robust results but require prohibitive manual efforts. Conversely, tetrahedral grids can be constructed more efficiently but their quality may be insufficient for traditional FEA due to poor aspect ratios. Overall, generation of quality meshes tends to consume significant resources, which creates a computational bottleneck in the FE workflow [14].

There are also situations such as Lagrangian simulations of large-deformation mechanics [19], where distorted grids are simply unavoidable. As a result, hardening finite element methods (FEM) against substandard grids can have significant impacts towards enabling automated CAD-to-solution capabilities by reducing or even removing the performance barriers created by the mesh-quality requirements of conventional FEA.

In this work we propose a new approach for the construction of FE shape functions that is independent of the underlying mesh. To that end we borrow ideas from meshfree methods and utilize Generalized Moving Least Squares (GMLS) [23] regression techniques to construct a FE basis whose degrees of freedom are not tied to the elements in the mesh. As a result, the role of the mesh in our approach is limited strictly to describing the computational domain and providing the means to integrate the basis functions without having any role in their approximation quality. Construction of our shape functions requires solution of small local quadratic programs on each element, and since the resulting basis is piecewise polynomial, integration of the polynomial products appearing in the assembly process can be accomplished with relatively few quadrature points.

Because our basis functions do not conform to the shape of the underlying elements, their span produces a *non-conforming*, discontinuous piecewise polynomial FE space. Thus, a discontinuous Galerkin (DG) [2, 11] and/or an Interior Penalty (IP) [24, 1] framework provide the most natural setting for the utilization of the new meshfree basis. This allows us to leverage the existing mature and rigorous stability and approximation theory for such methods by borrowing classical stabilization techniques from DG and IP to obtain robust and accurate FE discretizations that are also impervious to the quality of the underlying mesh.

The possibility to reuse DG and IP stabilization techniques and the fact that our basis functions only require integration of polynomial products sets the approach in this paper apart from other schemes such as meshfree Galerkin methods [5, 17, 3, 16], which replace traditional shape functions by meshfree ones. However, the meshfree bases in these methods are not polynomials and are not known in a closed form. As a result, their integration requires a relatively large number of quadrature points, which increases the computational cost of such schemes, as every shape function evaluation involves a solution of a small linear algebra problem. This has prompted consideration of reduced order integration [9, 8]; however, such integration leads to numerical instabilities due to underintegration and requires application-specific stabilizations.

We have organized the paper as follows. Section 2 reviews the technical background necessary for this work and includes a brief summary of the GMLS regression technique that is central to our approach. Section 3 explains the construction of the meshfree basis and demonstrates it use in a DG framework to obtain a well-posed partial differential equation (PDE) discretization. Further implementation details are discussed in Section 4 and numerical results are presented in Section 5. Our key findings and outlook for future work are summarized in Section 6. Collaborations and continuation of the products of this LDRD are described in Section 7.

2. TECHNICAL BACKGROUND

2.1. Notation

Throughout this paper Ω is a bounded open region in $\mathbb{R}^d$, $d = 1,2,3$ with a Lipschitz continuous boundary $\Gamma = \partial\Omega$. We assume that Γ comprises two disjoint parts Γ_D and Γ_N, respectively, i.e., $\Gamma = \Gamma_D \cup \Gamma_N$. We denote the outer unit normal to Γ by $\boldsymbol{n}$. As always, $L^2(\Omega)$ is the space of all square integrable functions on Ω, $H^k(\Omega)$ is the Sobolev space of order k, and $H^1_D(\Omega)$ is the subspace of $H^1(\Omega)$ containing all functions whose trace vanishes on Γ_D. We denote the norm and the inner product on these spaces by $\|\cdot\|_k$ and $(\cdot,\cdot)_k$ for $k = 0,1,\ldots$. For a nonempty $\gamma \in \Gamma$ the elements of $H^1(\Omega)$ have well defined traces in $H^{1/2}(\gamma)$. The norm and the inner product on this space are denoted by $\|\cdot\|_{1/2,\gamma}$, and $\langle\cdot,\rangle_{1/2,\gamma}$, respectively. We overload the symbol $|\dot{|}$ to denote a seminorm, measure or cardinality with the meaning determined by the type of its argument.

Let Ω_h denote a partition of Ω into finite elements $K_i \subset \mathbb{R}^d$ with diameters h_i, $i = 1,\ldots N$. We assume that Ω_h is conforming in the sense that its boundary Γ_h coincides with the boundary Γ of the underlying domain. Without a loss of generality we may also assume that $\Gamma_h = \Gamma_{h,D} \cup \Gamma_{h,N}$ with $\Gamma_{h,D} = \Gamma_D$ and $\Gamma_{h,N} = \Gamma_N$.

The set of all element interfaces is denoted by S_h, i.e.,

$$S_h = \cup_{K_i \in \Omega_h} \partial K_i \,,$$

while the set of all interior element interfaces is denoted by $\breve{S}_h$:

$$\breve{S}_h = S_h \setminus \Gamma_h$$

The mesh parameter is defined as

$$h = \max_{K_i \in \Omega_h} h_i \,.$$

Let $P_m(K_i)$ denote the space of all multivariate polynomials whose degree does not exceed some integer $m \geq 0$. We define a standard non-conforming finite element space on Ω_h as

$$V_h = \{ v_h \in L^2(\Omega) \,|\, v_h\big|_{K_i} \in P_m(K_i) \quad \forall K_i \in \Omega_h \} \,. \tag{2.1}$$

Every $\sigma \in \breve{S}_h$ is shared by exactly two elements, i.e., $\sigma = \partial K_i \cap \partial K_j$ for some indices i, j. Following the standard notation [13] we introduce the average and jump of a function w on σ as

$$\{w\}_\sigma := \frac{1}{2}\left(w\big|_{K_i} + w\big|_{K_j} \right) \quad \text{and} \quad [w]_\sigma := \left(w\big|_{K_i} - w\big|_{K_j} \right) ,$$

For $\sigma \in \Gamma_h$ we use the same notation with the understanding that

$$\{w\}_\sigma := w\big|_{K_i} \quad \text{and} \quad [w]_\sigma := w\big|_{K_i} \,,$$

where K_i is the unique element containing the segment σ.

2.2. Generalized Moving Least Squares

GMLS is a non-parametric regression technique for the approximation of bounded linear functionals from scattered data [23, Section 4.3]. A typical GMLS setting includes

- a function space U with a dual U^*;
- a finite dimensional space $\Phi \subset U$ with basis $\boldsymbol{\phi} = \{\phi_1, \ldots, \phi_q\}$;
- a Φ-unisolvent[1] set of sampling functionals $S = \{\lambda_1, \ldots, \lambda_n\} \subset U^*$; and
- a locally supported kernel $w : U^* \times U^* \mapsto \mathbb{R}^+ \cup \{0\}$.

GMLS seeks an approximation $\widetilde{\tau}(u)$ of the *target* $\tau(u) \in U^*$ in terms of the sample vector $\boldsymbol{u} := (\lambda_1(u), \ldots, \lambda_n(u)) \in \mathbb{R}^n$, such that $\widetilde{\tau}(\phi) = \tau(\phi)$ for all $\phi \in \Phi$, i.e., the approximation is Φ-reproducing. To define $\widetilde{\tau}(u)$ we need the vector $\boldsymbol{\tau}(\phi) \in \mathbb{R}^q$ with elements $(\boldsymbol{\tau}(\boldsymbol{\phi}))_i = \tau(\phi_i)$, $i = 1, \ldots, q$, the diagonal weight matrix $W(\tau) \in \mathbb{R}^{n \times n}$ with element $W_{ii}(\tau) = w(\tau; \lambda_i)$, and the basis sample matrix $B \in \mathbb{R}^{n \times q}$ with element $B_{ij} = \lambda_i(\phi_j)$; $i = 1, \ldots, n$; $j = 1, \ldots, q$. Let $|\cdot|_{W(\tau)}$ denote the Euclidean norm on $\mathbb{R}^n$ weighted by $W(\tau)$, i.e.,

$$|\boldsymbol{b}|^2_{W(\tau)} = \boldsymbol{b}^\intercal W(\tau) \boldsymbol{b} \qquad \forall \boldsymbol{b} \in \mathbb{R}^n .$$

The GMLS approximant of the target is then given by

$$\widetilde{\tau}(u) := \boldsymbol{c}(\boldsymbol{u}; \tau) \cdot \boldsymbol{\tau}(\phi), \tag{2.2}$$

where the GMLS coefficients $\boldsymbol{c}(\boldsymbol{u}; \tau) \in \mathbb{R}^q$ solve

$$\boldsymbol{c}(\boldsymbol{u}; \tau) = \operatorname*{argmin}_{\boldsymbol{b} \in \mathbb{R}^q} \frac{1}{2} |B\boldsymbol{b} - \boldsymbol{u}|^2_{W(\tau)} . \tag{2.3}$$

It is straightforward to check that $\boldsymbol{c}(\boldsymbol{u}; \tau) = (B^\intercal W(\tau) B)^{-1} (B^\intercal W(\tau)) \boldsymbol{u}$. We refer to [21] for information about the efficient and stable solution of (2.3). Lastly, let $\mathbf{e}_i \in \mathbb{R}^n$ be the ith Cartesian unit vector and let $u_i^\tau := \boldsymbol{c}(\mathbf{e}_i; \tau) \cdot \boldsymbol{\phi} \in \Phi$. We call the set $S^\tau = \{u_1^\tau, \ldots, u_n^\tau\} \subset \Phi$ *a GMLS reciprocal of S relative to τ.*

2.3. A framework for non-conforming Galerkin methods

An attractive feature of our approach is that the resulting non-conforming finite element space can be used in conjunction with any of the available non-conforming discretization frameworks such as DG and IP. In this work we use extensively a setting motivated by symmetric interior penalty Galerkin (SIPG) methods as described in [13]. For completeness we review the fundamentals of this setting below.

[1] We recall that Φ-unisolvency implies $\{\phi \in \Phi \,|\, \lambda_i(\phi) = 0, i = 1, \ldots, n\} = \{0\}$.

As a model problem we consider the following elliptic reaction-diffusion equation

$$\begin{cases} -\nabla\cdot(\kappa\nabla u)+\alpha u &= f \quad \text{in } \Omega \\ u &= u_D \quad \text{on } \Gamma_D \\ \kappa\nabla u\cdot\boldsymbol{n} &= u_N \quad \text{on } \Gamma_N \end{cases} \tag{2.4}$$

where $f\in L^2(\Omega)$, $u_D\in H^{1/2}(\Gamma_D)$, and $u_N\in L^2(\Gamma_N)$ are given functions, α is a nonnegative scalar function, and κ is a $d\times d$ symmetric and positive definite matrix function.

For simplicity we shall assume that Γ_D is non-empty, which ensures that (2.4) has a unique solution. A standard SIPG weak formulation of (2.4) is given by: *seek $u_h\in V_h$ such that*

$$Q(u_h,v_h)=F(v_h)\quad \forall v_h\in V_h \tag{2.5}$$

where the bilinear form $Q:V_h\times V_h\mapsto\mathbb{R}$ and the linear functional $F:V_h\mapsto\mathbb{R}$ are defined as

$$\begin{aligned} Q(w_h,v_h)=&\sum_{K_i\in\Omega_h}\int_{K_i}\kappa\nabla w_h\cdot\nabla v_h dx+\int_\Omega \alpha w_h v_h dx+\sum_{\sigma\in\Gamma_{h,D}}\tau(\sigma)\int_\sigma [w_h][v_h]d\ell \\ &-\sum_{\sigma\in\Gamma_{h,D}}\int_\sigma\{\kappa\nabla w_h\cdot\boldsymbol{n}\}[v_h]d\ell-\sum_{\sigma\in\Gamma_{h,D}}\int_\sigma\{\kappa\nabla v_h\cdot\boldsymbol{n}\}[w_h]d\ell, \end{aligned} \tag{2.6}$$

and

$$\begin{aligned} F(v_h)=&\int_\Omega f v_h dx-\sum_{\sigma\in\Gamma_{h,D}}\int_\sigma(\kappa\nabla v_h\cdot\boldsymbol{n})u_D d\ell \\ &+\sum_{\sigma\in\Gamma_{h,D}}\tau(\sigma)\int_\sigma u_D v_h d\ell+\sum_{\sigma\in\Gamma_{h,N}}\int_\sigma v_h u_N d\ell, \end{aligned} \tag{2.7}$$

respectively. Here, $\tau(\sigma)$ is a stabilization (penalty) parameter for the interface σ. For SIPG methods defined with standard non-conforming spaces this parameter is commonly defined as

$$\tau(\sigma):=\frac{\varepsilon_\sigma}{|\sigma|^{\beta_0}}$$

where ε_σ is a positive constant on each interface and β_0 is a global constant usually set to $(d-1)^{-1}$; see [13].

The bilinear form (2.6) induces an "energy" norm

$$\|v_h\|_E=\left(\sum_{K_i\in\Omega_h}\int_{K_i}\kappa|\nabla v_h|^2dx+\int_\Omega\alpha v_h^2dx+\sum_{\sigma\in\Gamma_{h,D}}\tau(\sigma)\int_\sigma[v_h]^2d\ell\right) \tag{2.8}$$

Under some conditions one can show that $Q(\cdot,\cdot)$ is continuous and coercive with respect to the energy norm (2.8); see [13] and the references therein, and prove apriori error estimates with respect to this norm.

In Section 3 we will use this SIPG framework in conjunction with our meshfree basis functions. We wills see that, among other things, the use of the meshfree basis allows to relax the definition of the penalty parameter $\tau(\sigma)$.

3. A NONCONFORMING VARIATIONAL FORMULATION

In this section we use the GMLS regression technique to define a finite element basis independently of the underlying mesh. At the end of the section we briefly explain how our approach can be viewed as an extension of GMLS to the approximation of bilinear forms.

3.1. A nonconforming finite element space based on GMLS regression

Given a finite element partition Ω_h we assume that $X_h = \{\boldsymbol{x}_i\}_{i=1}^R$ is a point cloud on Ω defined independently of Ω_h. We note that X_h may contain some or all of the element vertices but in general it does not have to be dependent on Ω_h in any way.

To define our new basis we specialize the general GMLS framework in Section 2.2 as follows. For simplicity we set $U = U^* = C^0(\Omega)$ and define $\Phi = P_m(\mathbb{R}^d)$, i.e., the space of all multivariate polynomials in d-dimensions. To avoid notational clutter we retain the notation $\boldsymbol{\phi}$ for the basis of P_m. Then we choose S to be the set of all point samples at the cloud points, i.e., we set $\lambda_i = \delta(\boldsymbol{x}_i)$, where δ is the Dirac delta function. Thus,

$$S = \{\delta(\boldsymbol{x}_1), \ldots, \delta_{\boldsymbol{x}_R}\}. \tag{3.1}$$

We assume that the points comprising X_h satisfy the unisolvency requirements in [23].

Let $\boldsymbol{b}_k$ denote the barycenter of element K_i. We define the set of target functionals τ as

$$T = \{\delta(\boldsymbol{b}_1), \ldots, \delta(\boldsymbol{b}_N)\}.$$

Finally, we set the kernel w as

$$w(\mathscr{K}_k, \boldsymbol{x}_j) := \rho(|\boldsymbol{b}_k - \boldsymbol{x}_j|),$$

where $\rho(\cdot)$ is a radially symmetric function with $\operatorname{supp}\rho = O(h)$. The GMLS basis will be constructed locally from point samples close to $\boldsymbol{b}_k$ using the local sampling set $S_k = \{\delta_{\boldsymbol{x}_j} \,|\, w(\mathscr{K}_k, \boldsymbol{x}_j) > 0\}$ with cardinality n_k. We assume that the support of ρ is large enough to ensure that S_k is P^m-unisolvent.

We then define the (local) finite element GMLS basis as the GMLS reciprocal of the delta function sampling set (3.1). Evaluation of this basis depends on the selection of the target functional and so we have that

$$S^k = \{u_1^k, \ldots, u_{n_k}^k\}$$

with $u_i^k := \boldsymbol{c}(\mathbf{e}_i^k; \boldsymbol{b}_k) \cdot \boldsymbol{\phi}$ and $\mathbf{e}_i^k \in \mathbb{R}^{n_k}$.

We now use this basis in conjunction with the SIPG formulation (2.5) of the model problem (2.4). To that end, let Q_k and F_k denote the restrictions of the SIPG bilinear form and right hand side

functional to an element $K_k \in \Omega_h$. Further restricting Q_k to $S^k \times S^k$ yields the elemental assembly contribution

$$Q_k(u_j^k, u_i^k) := Q_k(\boldsymbol{c}(\mathbf{e}_j^k; \boldsymbol{b}_k)\boldsymbol{\phi}, \boldsymbol{c}(\mathbf{e}_i^k; \boldsymbol{b}_k)\boldsymbol{\phi}) = \boldsymbol{c}(\mathbf{e}_i^k; \boldsymbol{b}_k) \cdot a_k(\boldsymbol{\phi}, \boldsymbol{\phi}) \cdot \boldsymbol{c}(\mathbf{e}_j^k; \boldsymbol{b}_k)$$

from element K_k. Similarly, we have the elemental approximation

$$F_k(u_i^k) := F_k(\boldsymbol{c}(\mathbf{e}_i^k; \boldsymbol{b}_k)\boldsymbol{\phi}) = \boldsymbol{c}(\mathbf{e}_i^k; \boldsymbol{b}_k) \cdot F_k(\boldsymbol{\phi})$$

of the right hand side functional. It is easy to see that the elemental approximations give rise to a local assembly matrix

$$\mathbf{Q}_{ij}^k = \boldsymbol{c}(\mathbf{e}_i^k; \boldsymbol{b}_k) \cdot Q_k(\boldsymbol{\phi}, \boldsymbol{\phi}) \cdot \boldsymbol{c}(\mathbf{e}_j^k; \boldsymbol{b}_k) \tag{3.2}$$

and a local load vector $\mathbf{F}_i^k = \boldsymbol{c}^v(\mathbf{e}_i; \tau) \cdot F(\boldsymbol{\phi})$, respectively, which are analogues of the element stiffness matrix and load vector in FEA. With this local problem we associate the local set of degrees of freedom $\boldsymbol{a}^k$.

It is important to note that unlike some meshfree Galerkin methods, assembly of the local SIPG element contributions only involves integration of polynomial functions. Specifically, computation of element contributions boils down to computing the matrix $Q_k(\boldsymbol{\phi}, \boldsymbol{\phi}) \in \mathbb{R}^{n_q \times n_q}$ with element

$$(Q_k(\boldsymbol{\phi}, \boldsymbol{\phi}))_{st} = Q_k(\boldsymbol{\phi}_s, \boldsymbol{\phi}_t)$$

and the vector $F_k(\boldsymbol{\phi}) \in \mathbb{R}^{n_q}$ with element

$$(F_k(\boldsymbol{\phi}))_s = F_k(\phi_s)\,.$$

Since ϕ_s and ϕ_t are polynomials, computation of the necessary integrals can be done by standard numerical quadrature. In some cases these integrals can also be computed exactly.

Let us now briefly discuss the global finite element space corresponding to the local spaces S^k. We construct this space as $[S] = \cup_{K_k \in \Omega^h} S^k$ and denote its elements by $[u]$. Stacking all local DoF in a single vector $[\boldsymbol{a}] := \{\boldsymbol{a}^1, \ldots, \boldsymbol{a}^{N_e}\}$ produces the global DoF set for $[u]$. In general, a sampling functional $\delta_{\boldsymbol{x}_j}$ can belong to multiple local sampling sets S_k, which means that $[u]$ will be multivalued at $\boldsymbol{x}_j$, i.e., the global space $[S]$ is discontinuous along the element interfaces.

We now define the global SIPG problem by summing over all elements, i.e.,

$$Q([u],[v]) := \sum_{K_k \in \Omega^h} Q_k(u^k, v^k) \quad \text{and} \quad F([v]) := \sum_{K_k \in \Omega^h} F_k(v^k)\,,$$

where $u^k, v^k \in S^k$. Implementation details are further discussed in Section 4.

3.2. Interpretation as extension of GMLS to bilinear forms

As mentioned in Section 2.2, GMLS is a non-parametric regression technique for the approximation of bounded linear functionals. Below we show that our approach can be interpreted as an extension of GMLS to the approximation of the SIPG bilinear form $Q(\cdot,\cdot)$.

Because this interpretation applies to a broad class of variational formulations, we sketch the details using the abstract setting of Section 2.2 rather than the specialized setting from Section 3.1. Thus we consider a bilinear form $Q : U \times V \to \mathbb{R}$, where U and V denote Hilbert spaces with duals U^* and V^* and a generic variational equation *given* $F \in V^*$ *find* $u \in U$ *such that*

$$Q(u,v) = F(v) \qquad \forall v \in V\,. \tag{3.3}$$

Following the usual practice we refer to U and V as the trial and the test space, respectively.

The key observation that enables the extension of GMLS from linear functionals to bilinear forms is the fact that for a fixed trial function $u \in U$ the form $Q(u,\cdot)$ defines a bounded linear functional on V, i.e., $Q(u,\cdot) \in V^*$, while for a fixed test function $v \in V$ we have that $Q(\cdot,v) \in U^*$. The details follow.

To approximate (3.3) we will use two separate instances of the GMLS regression for the test and trial spaces, respectively. To differentiate between these instances we tag their entities with a sub/superscript indicating the underlying space, e.g., S_U and $\boldsymbol{c}^U(u)$ denote a sampling set and a coefficient vector, respectively, for the *trial space*. One exception to this rule will be the reciprocal GMLS functions u_i^τ and v_i^τ.

We obtain the GMLS approximation of (3.3) in two steps. As already explained, for any fixed $u \in U$ the form $Q(u,\cdot)$ defines a bounded linear functional on V, i.e., $Q(u,\cdot) \in V^*$. We shall assume that the kernel w is such that $W(a(u,\cdot)) = W(f)$. For this reason we retain the generic label τ to indicate dependence of various GMLS entities on their respective target functionals. Then, the GMLS approximants of $Q(u,\cdot)$ and f can be written in terms of the same GMLS coefficient vector as

$$\widetilde{Q}(u,v) := \boldsymbol{c}^V(v;\tau)\cdot Q(u,\boldsymbol{\phi}^V) \quad \text{and} \quad \widetilde{F}(v) = \boldsymbol{c}^V(v;\tau)\cdot F(\boldsymbol{\phi}^V) \quad \forall v \in V\,,$$

respectively. Combining these representations yields the following approximation of (3.3): *find* $u \in U$ *such that* $\widetilde{Q}(u,v) = \widetilde{F}(v)$ for all $v \in V$, or equivalently,

$$\boldsymbol{c}^V(v;\tau)\cdot Q(u,\boldsymbol{\phi}^V) = \boldsymbol{c}^V(v;\tau)\cdot F(\boldsymbol{\phi}^V) \qquad \forall v \in V. \tag{3.4}$$

The weak problem (3.4) has infinitely many "equations" and "variables". To reduce the number of equations we restrict the test space in (3.4) to the GMLS reciprocal set S_V^τ to obtain the following problem: *find* $u \in U$ *such that* $\widetilde{Q}(u,v_i^\tau) = \widetilde{F}(v_i^\tau)$ *for all* $v_i^\tau \in S_V^\tau$ or, which is the same,

$$\boldsymbol{c}^V(\mathbf{e}_i;\tau)\cdot Q(u,\boldsymbol{\phi}^V) = \boldsymbol{c}^V(\mathbf{e}_i;\tau)\cdot F(\boldsymbol{\phi}^V) \qquad i = 1,\ldots,n^V\,. \tag{3.5}$$

This completes the first step. The second step discretizes the trial space by restricting the search for a solution to the reciprocal GMLS space S_U^τ, i.e., we consider the problem: *find* $u^\tau \in S_U^\tau$ *such that*

$$\widetilde{Q}(u^\tau,v_i^\tau) = \widetilde{F}(v_i^\tau) \qquad \forall v_i^\tau \in S_V^\tau. \tag{3.6}$$

where $u^\tau := \sum_{j=1}^{n^U} a_j u_j^\tau$. Using (3.5) and $u_j^\tau := \boldsymbol{c}(\mathbf{e}_j;\tau)\cdot\boldsymbol{\phi}$ one can write (3.6) as

$$\sum_{j=1}^{n^U} \left(\boldsymbol{c}^V(\mathbf{e}_i;\tau)\cdot Q(\boldsymbol{\phi}^U,\boldsymbol{\phi}^V)\cdot\boldsymbol{c}^U(\mathbf{e}_j;\tau)\right) a_j = \boldsymbol{c}^V(\mathbf{e}_i;\tau)\cdot F(\boldsymbol{\phi}^V) \quad i = 1,\ldots,n^V\,. \tag{3.7}$$

It is easy to see that this problem is equivalent to the following $n^V \times n^U$ system of linear algebraic equations

$$\mathbf{Q}\boldsymbol{a} = \mathbf{F} \tag{3.8}$$

where $\boldsymbol{a} = \{a_1, \ldots, a_{n^U}\} \in \mathbb{R}^{n^U}$ are the GMLS degrees-of-freedom (DoF), while $\mathbf{Q} \in \mathbb{R}^{n^V \times n^U}$ and $\mathbf{F} \in \mathbb{R}^{n^V}$ have elements

$$Q_{ij} = \boldsymbol{c}^V(\mathbf{e}_i;\tau) \cdot Q(\boldsymbol{\phi}^U, \boldsymbol{\phi}^V) \cdot \boldsymbol{c}^U(\mathbf{e}_j;\tau) \quad \text{and} \quad F_i = \boldsymbol{c}^V(\mathbf{e}_i;\tau) \cdot F(\boldsymbol{\phi}^V),$$

respectively. Problems (3.6) and (3.8) can be viewed as GMLS analogues of variational discretization of (3.3) and its equivalent linear algebraic form, respectively. It is clear, however, that the algebraic problem resulting from this perspective is completely equivalent to the algebraic problem involving (3.2), which was obtained by treating the reciprocal fields u_j^τ and v_i^τ as "finite element basis functions" for the trial and test spaces. We refer to [6] for additional details of this interpretation of our approach.

4. IMPLEMENTATION DETAILS

We use an in-house C++ code that joins together components found in Trilinos [22] for managing parameter lists (Teuchos), file input and output, creating and managing fields, creating graphs and matrices for storing the linear systems (Tpetra), solving the linear systems (Amesos2, a direct solver), generating quadrature points and weights (Intrepid), and performing the other routines common to a DG-FEM simulation code.

The quadratic programs posed by GMLS are assembled and solved with the Compadre Toolkit [15], a package in Trilinos for portable and parallel (multi-threaded and can run on GPUs by using Kokkos [12]) creation of dense matrices relating to GMLS and computing their solution. This toolkit also exists in stand-alone form and can be easily added to existing DG-FEM codes.

As mentioned previously, one GMLS problem must be solved per cell in the mesh. This is substantially less work than having to solve at each quadrature point. The "neighbors" needed for GMLS are the modal degrees of freedom from the particle set.

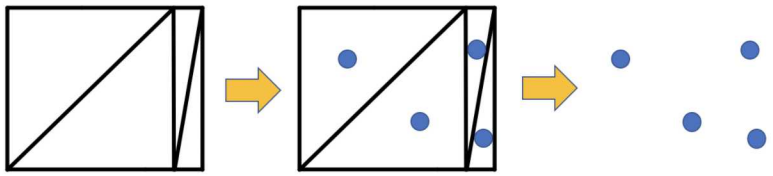

Figure 4-1. Particle set generated from centroids of a mesh

There are an infinite number of choices for the degrees of freedom (the particles), but we choose to consider three options here. The first option is the easiest to implement, namely using the centroids of the cells on the mesh as the particle cloud of degrees of freedom. If a mesh is highly distorted however, cells sizes in certain regions of the domain may be much smaller than cell sizes in other regions of the domain, such that their centroids as a particle cloud may have extreme clustering and result in a poor quality GMLS solution. This type of particle set and its generation is shown in Figure 4-1.

The second option we consider for choice of particle locations is a uniform to quasi-uniform filling of the geometry with points. Specifying a different particle set than the centroids of the cells of the mesh under consideration requires additional software considerations, but fits fully within the generality of our framework. The generation of such particle sets can be achieved by a uniform filling of a bounding box in Cartesian coordinates that then have particles not inside the geometry removed. This type of particle set and its generation is shown in Figure 4-2. Here the mesh outline is shown simply to indicate that it is disregarded in the point generation.

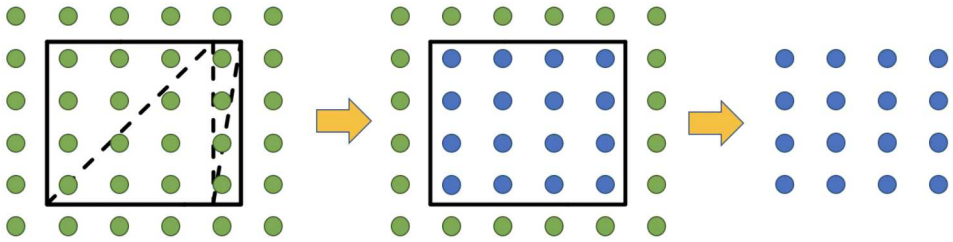

Figure 4-2. Particle set generated from Cartesian bounding box

A third option, relevant to the scenario where a mesh has become internally highly-deformed under some process, is to use the centroids from some other mesh having more uniform centroids as a particle set. For instance, in large deformation solid mechanics one typically has access to both a reference and physical configuration corresponding to the configuration of the material before and after loading. This assumes some reference/initial configuration exists and has cell centroids that are at least quasi-uniformly spaced. Options one and three are the two particle set types that we will choose from in numerical studies.

5. NUMERICAL RESULTS

We begin our numerical investigation by describing a sequence of pathological mesh configurations that will clearly demonstrate the problems faced by traditional FEM under high distortion.

5.1. Highly distorted mesh sequences

Instead of defining the geometry of an annulus directly, we discretize a rectangle of shape $(r,\theta) \in [0.5, 1.0] \times [0, 2\pi]$. For the initial discretization (mesh level 0), we use 6 elements horizontally and 24 elements vertically. For each successive level of refinement, the number of elements is doubled both horizontally and vertically. Following the generation of the rectangular meshes just defined, the nodes are moved under the transformation:

$$\begin{bmatrix} x \\ y \end{bmatrix} = \begin{bmatrix} r\,cos(\theta + tf\,(r/(r_o - r_i) - r_o/(r_o - r_i))) \\ r\,sin(\theta + tf\,(r/(r_o - r_i) - r_o/(r_o - r_i))) \end{bmatrix}$$

where r_o is the outer radius of annulus, r_i is the inner radius of the annulus, and tf is the torsion factor. We consider three annuli of varying tf values, which we vary from 1.0 (well-conditioned elements) to 1.795 (near element inversion).

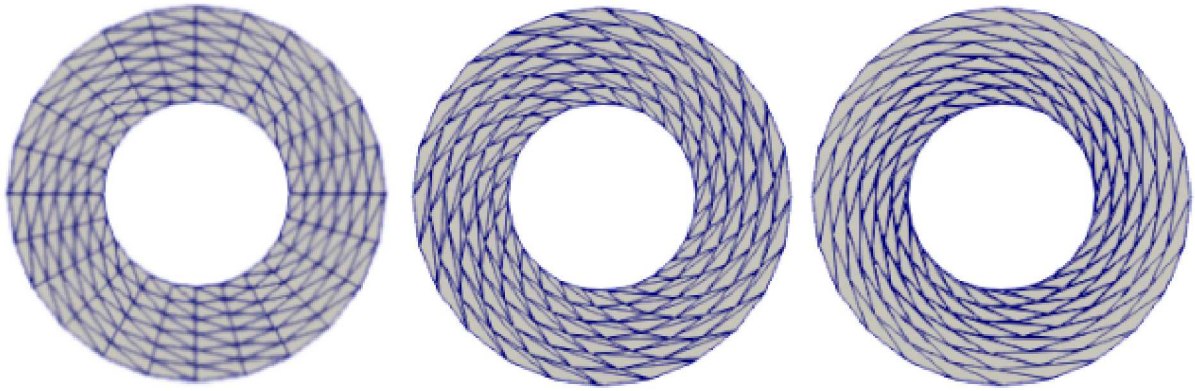

Figure 5-1. Varying degrees of element shear for a.) $tf = 0$ (no shear, reference[left]), b.) $tf = 1.0$ (moderate shear, shear level 1[center]), and c.) $tf = 1.795$ (extreme shear, shear level 2[right])

In Figure 5-2, we zoom into Figure 5-1(right) to observe the extremely large angles (nearly π) and elements having area approaching zero. In addition to having elements with an angle

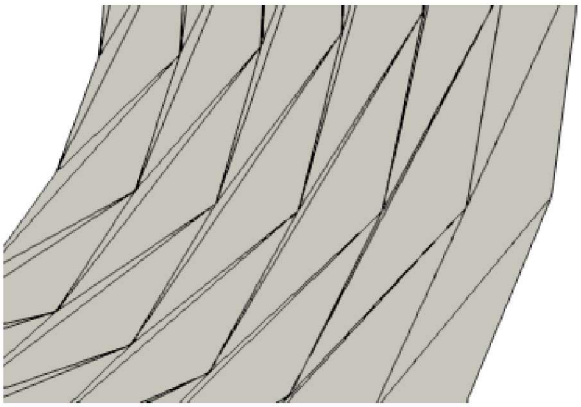

Figure 5-2. Zoom-in to Figure 5-1c

approaching π, notice that the elements are of increasing size when approaching the outer portion of the annulus from the inside.

An anticipated difficulty in using the centroids of these cells for the choice of particles is that the extremely sheared annulus has elements whose centroids are clustered towards the inner annulus and spread out towards the outer annulus. Specifying a uniform window size (to ensure symmetry of neighbor lists) for each GMLS problem requires a larger window size to accommodate the sparse placement of centroids towards the outer annulus.

Centroids of the non-sheared mesh are essentially uniformly spaced, enabling a smaller uniform window size to be used for the GMLS problems if selected as the particle set. This should yield increased accuracy when performing computations on the extremely sheared annulus in tandem with particle sets defined on the centroids of the non-sheared mesh.

In the section on reaction-diffusion to follow, we will investigate the properties listed for both particle set choices, referring to particle locations chosen as the centroids of the mesh as "centroid particles" and particle locations chosen from a reference or non-sheared annulus as "reference particles." We break them into subsections based on whether the particle sets are derived from the mesh being used and using the window size for GMLS set to accommodate the most deformed mesh in that subsection, or whether the particle sets are derived from the no-shear annulus centroids, used as a reference configuration, with a window size for GMLS set to accommodate the no-shear meshes. There will be a final subsection that will compare and contrast the differences in behavior between the two choices of particle sets.

The approach outlined in this document is sufficiently general as to cover scalar and vector elliptic PDEs as well as incompressible linear elasticity and Stokes flow, so we now turn our attention toward numerical studies of each application. For each application, we will describe problems often exhibited for traditional FEs, and then will present a comparison of quantities relevant to the stated challenge, comparing the GMLS basis function + SIPG framework to traditional nodal continuous Galerkin (CG) FEM.

5.2. Strongly coercive scalar elliptic problems

For this first case, we present results measuring rates of convergence for scalar reaction-diffusion problems posed on the sheared annuli meshes and also compare against traditional CG-FEM solutions on the same meshes. We use the method of manufactured solutions to generate boundary conditions and body forces consistent with reaction parameter $\alpha = 1$, diffusion parameter $\kappa = 1$, and an exact solution $(\sin(x) \cdot \sin(y))$. We solve the PDE and then compare the computed solution to the exact solution.

The SIPG framework admits flexibility in the choice of σ_e as well as β, the power for the edge length, $|e|$. An additional length $|e|$ that can be chosen is the neighbor search size (window size). In Tables 5-1-5-3 and Tables 5-4-5-6, we use a global choice for the constant $\sigma_e = 100$, $\beta = 1$, and vary the choice of $|e|$ from 1, cell edge length (traditional SIPG), and neighbor search size. The numerical results indicate that all three choices for $|e|$ provide optimally convergent results.

In the tables contained in the following two sections, errors and rates are given in the L^2 norm and *SIPG* norm (2.8).

5.2.1. Particles from mesh centroids

Due to large disparity in element sizes for the extreme-shear annulus meshes, a neighbor radius / window size of 0.55 is used for all numerical results that follow, even though a radius half that size is sufficiently large to ensure optimal convergence of the no-shear annulus case.

Table 5-1. Convergence study for tf=0.0
(annulus with no shear using mesh particles)

L^2 Norm						
Refinement	$\|e\|$=1		$\|e\|$=window		$\|e\|$=edge length	
Level	Error	Rate	Error	Rate	Error	Rate
0	2.9819e-04	-	4.5970e-04	-	1.1489e-03	-
1	8.2033e-06	5.18	1.9589e-05	4.55	5.3945e-05	4.41
2	2.4382e-07	5.07	5.2023e-07	5.23	1.2719e-06	5.41
3	1.4044e-08	4.12	2.3853e-08	4.45	4.4603e-08	4.83
4	8.2321e-10	4.09	1.1532e-09	4.37	1.8584e-09	4.59
SIPG Norm						
Refinement	$\|e\|$=1		$\|e\|$=window		$\|e\|$=edge length	
Level	Error	Rate	Error	Rate	Error	Rate
0	1.5618e-02	-	2.0491e-02	-	3.3359e-02	-
1	1.3119e-03	3.57	2.3946e-03	3.10	4.3971e-03	2.92
2	1.4734e-04	3.15	3.8378e-04	2.64	6.5204e-04	2.75
3	1.4316e-05	3.36	5.1265e-05	2.90	8.4994e-05	2.93
4	1.3213e-06	3.44	6.4844e-06	2.98	1.0640e-05	3.00

Using 3^{rd} order polynomial basis in GMLS.

Table 5-2. Convergence study for tf=1.0
(annulus with moderate shear using mesh particles)

L^2 Norm						
Refinement	$\|e\|$=1		$\|e\|$=window		$\|e\|$=edge length	
Level	Error	Rate	Error	Rate	Error	Rate
0	6.3877e-04	-	9.2152e-04	-	1.6844e-03	-
1	1.6751e-05	5.25	4.6649e-05	4.30	1.0994e-04	3.94
2	4.8592e-07	5.11	1.5910e-06	4.87	4.1187e-06	4.74
3	2.4832e-08	4.29	4.6319e-08	5.10	1.0202e-07	5.34
4	1.5856e-09	3.97	2.1131e-09	4.45	3.5482e-09	4.85
SIPG Norm						
Refinement	$\|e\|$=1		$\|e\|$=window		$\|e\|$=edge length	
Level	Error	Rate	Error	Rate	Error	Rate
0	2.9959-e02	-	3.9475e-02	-	6.2601e-02	-
1	2.8643e-03	3.39	5.3201e-03	2.89	8.7145e-03	2.84
2	2.6944e-04	3.41	7.1298e-04	2.90	1.1756e-03	2.89
3	2.4038e-05	3.49	8.8786e-05	3.01	1.4631e-04	3.01
4	2.2493e-06	3.42	1.1090e-05	3.00	1.8289e-05	3.00

Using 3^{rd} order polynomial basis in GMLS.

Table 5-3. Convergence study for tf=1.795
(annulus with extreme shear using mesh particles)

L^2 Norm						
Refinement	$\|e\|$=1		$\|e\|$=window		$\|e\|$=edge length	
Level	Error	Rate	Error	Rate	Error	Rate
0	9.0164e-04	-	1.2038e-03	-	1.6369e-03	-
1	3.7529e-05	4.59	1.0055e-04	3.58	1.7433e-04	3.23
2	9.7529e-07	5.27	4.0551e-06	4.63	8.1635e-06	4.42
3	3.9421e-08	4.63	1.2517e-07	5.02	2.3224e-07	5.14
4	2.4772e-09	3.99	3.7143e-09	5.07	5.8397e-09	5.31
SIPG Norm						
Refinement	$\|e\|$=1		$\|e\|$=window		$\|e\|$=edge length	
Level	Error	Rate	Error	Rate	Error	Rate
0	4.7235e-02	-	6.2603e-02	-	9.3540e-02	-
1	4.8339e-03	3.29	8.9325e-03	2.81	1.3065e-02	2.84
2	4.3058e-04	3.49	1.1383e-03	2.97	1.6958e-03	2.95
3	3.8513e-05	3.48	1.4358e-04	2.99	2.1504e-04	2.98
4	3.6014e-06	3.42	1.8024e-05	2.99	2.6985e-05	2.99

Using 3^{rd} order polynomial basis in GMLS.

5.2.2. Particles from undeformed reference mesh centroids

Using the non-sheared annulus reference mesh centroids as the particle set, an initial neighbor radius / window size of 0.275 is sufficient to capture enough neighbors for the GMLS problem to be well posed at each centroid on the deformed meshes. Due to this decreased stencil size needed in tandem with the more uniform particle set, the errors in this subsection are improved relative to the errors in Section 5.2.

Table 5-4. Convergence study for tf=0.0 (annulus with no shear using reference particles)

L^2 Norm						
Refinement	$\|e\|$=1		$\|e\|$=window		$\|e\|$=edge length	
Level	Error	Rate	Error	Rate	Error	Rate
0	2.8420e-05	-	5.9814e-05	-	9.6072e-05	-
1	1.3693e-06	4.38	4.4807e-06	3.74	5.5028e-06	4.13
2	6.3880e-08	4.42	2.3879e-07	4.23	2.6866e-07	4.36
3	4.4490e-09	3.84	1.0360e-08	4.53	1.1275e-08	4.57
4	2.1703e-10	4.36	4.4678e-10	4.54	4.7909e-10	4.56
SIPG Norm						
Refinement	$\|e\|$=1		$\|e\|$=window		$\|e\|$=edge length	
Level	Error	Rate	Error	Rate	Error	Rate
0	3.2063e-03	-	5.9504e-03	-	7.6753e-03	-
1	4.2330e-04	2.92	1.1094e-03	2.42	1.3573e-03	2.50
2	4.2833e-05	3.30	1.5794e-04	2.81	1.8540e-04	2.87
3	4.3602e-06	3.30	2.0658e-05	2.93	2.3888e-05	2.96
4	3.6354e-07	3.58	2.6161e-06	2.98	3.0078e-06	2.99

Using 3^{rd} order polynomial basis in GMLS.

Table 5-5. Convergence study for tf=1.0
(annulus with moderate shear using reference particles)

L^2 Norm						
Refinement	$\|e\|$=1		$\|e\|$=window		$\|e\|$=edge length	
Level	Error	Rate	Error	Rate	Error	Rate
0	3.0290e-05	-	6.0610e-05	-	8.9921e-05	-
1	1.2340e-06	4.62	3.6212e-06	4.07	4.6223e-06	4.28
2	6.7731e-08	4.19	1.6355e-07	4.47	1.9528e-07	4.57
3	4.2708e-09	3.99	9.9464e-09	4.04	1.0738e-08	4.18
4	2.3526e-10	4.18	5.3849e-10	4.21	5.7320e-10	4.23
SIPG Norm						
Refinement	$\|e\|$=1		$\|e\|$=window		$\|e\|$=edge length	
Level	Error	Rate	Error	Rate	Error	Rate
0	4.8296e-03	-	9.0950e-03	-	1.0629e-02	-
1	3.7377e-04	3.69	9.8365e-04	3.21	1.1458e-03	3.21
2	3.4075e-05	3.46	1.2502e-04	2.98	1.4562e-04	2.98
3	3.3937e-06	3.33	1.5991e-05	2.97	1.8136e-05	3.01
4	2.9528e-07	3.52	1.9869e-06	3.01	2.2505e-06	3.01

Using 3^{rd} order polynomial basis in GMLS.

Table 5-6. Convergence study for tf=1.795
(annulus with extreme shear using reference particles)

L^2 Norm						
Refinement	$\|e\|$=1		$\|e\|$=window		$\|e\|$=edge length	
Level	Error	Rate	Error	Rate	Error	Rate
0	7.1015e-05	-	1.7989e-04	-	1.7919e-04	-
1	1.9263e-06	5.204	7.7180e-06	4.542	9.0628e-06	4.305
2	8.5845e-08	4.487	2.8712e-07	4.748	3.2912e-07	4.783
3	5.3423e-09	4.006	2.0561e-08	3.803	2.2504e-08	3.870
4	2.7279e-10	4.291	1.4786e-09	3.797	1.6141e-09	3.801
SIPG Norm						
Refinement	$\|e\|$=1		$\|e\|$=window		$\|e\|$=edge length	
Level	Error	Rate	Error	Rate	Error	Rate
0	5.7678e-03	-	1.0640e-02	-	1.1280e-02	-
1	5.3163e-04	3.439	1.3949e-03	2.931	1.4450e-03	2.964
2	4.8434e-05	3.456	1.7766e-04	2.972	1.8368e-04	2.975
3	4.7910e-06	3.337	2.3325e-05	2.929	2.3959e-05	2.938
4	4.2448e-07	3.496	2.9106e-06	3.002	2.9563e-06	3.018

Using 3^{rd} order polynomial basis in GMLS.

5.2.3. Particle set choice effects

Choice of $|\mathbf{e}|$ It is noteworthy that in either choice of particle sets, a choice of $|e| = 1$ works to stabilize this approach, because this means that the penalization for the jump terms may not need to increase with mesh refinement [Cf. Tables 5-1-5-3 and Tables 5-4-5-6]. While not particularly important for reaction-diffusion, this does have an effect on the largest stable timestep size for linear elasticity, which will be shown in the next section.

Comparison with CG-FEM Scalar elliptic problems, e.g. conjugate-heat transfer, are not known for being as significantly affected by poor mesh quality as problems in the solid mechanics community, such as elasticity and plasticity. For that reason, we do not expect the GMLS+SIPG approach to do significantly better than CG-FEM.

Figures 5-3 and 5-4 compare the GMLS+SIPG approach for both choices of particle sets against traditional CG-FEM. In the L^2 norm and H^1 seminorm, our approach is able to beat the solution provided by CG-FEM for the most extremely deformed mesh when the reference particles are used. Also note that using the reference particle set results in a nearly-flat error across all levels of annulus shear. Even for the mesh centroid particle set, the error rises at a slower rate than CG-FEM, but begins significantly worse-off on the undeformed annulus mesh.

Given that CG-FEM is expected to give decent solutions even for poor quality meshes, the ability of our approach to give better results on poor quality meshes is significant, and will be made especially evident in the next section on linear elasticity.

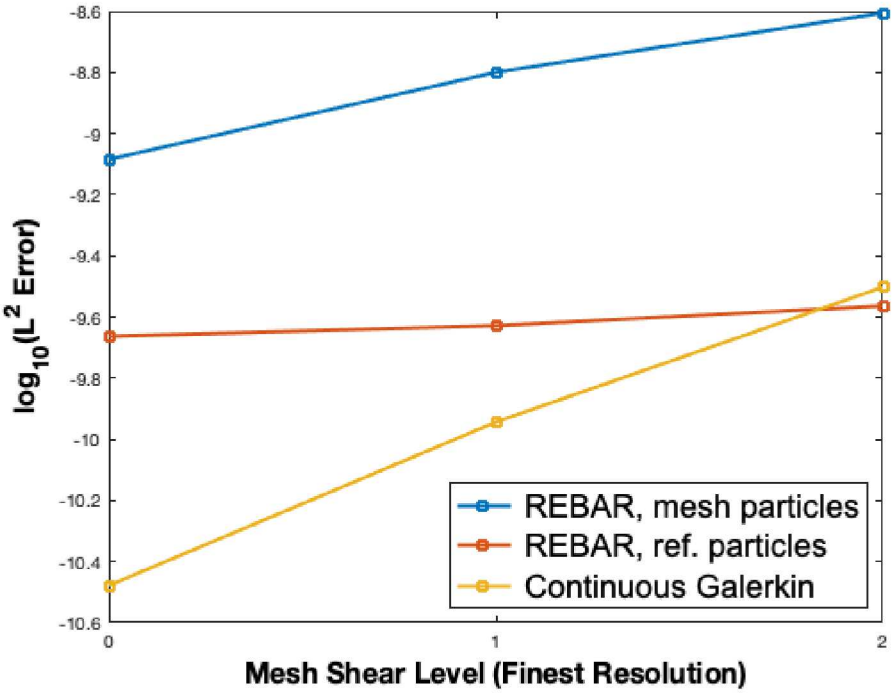

Figure 5-3. Comparison of L^2 norm of GMLS+SIPG against CG using a third order polynomial basis

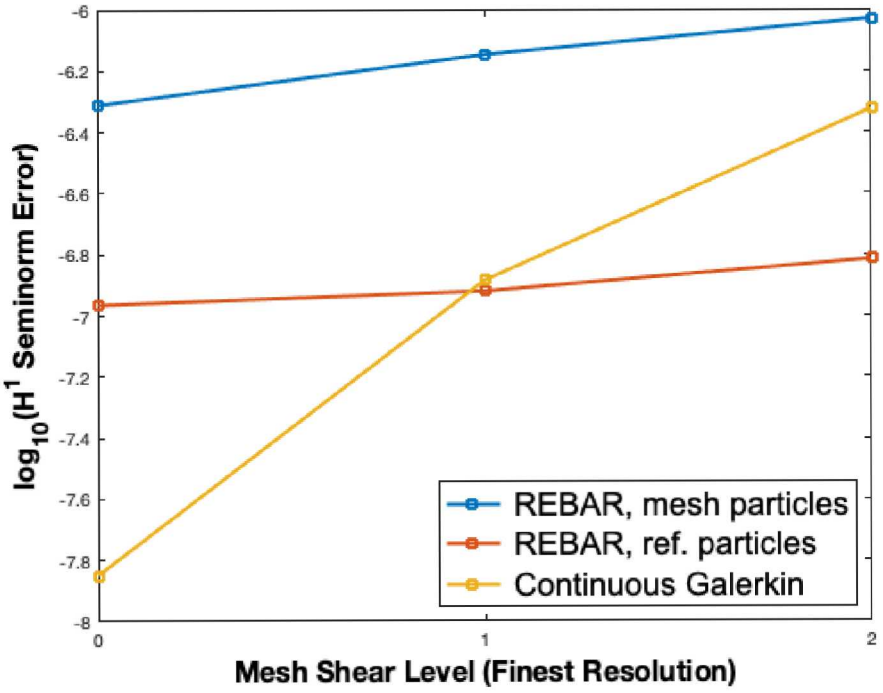

Figure 5-4. Comparison of H^1 seminorm of GMLS+SIPG against CG using a third order polynomial basis

5.3. Linear elasticity

The detrimental effects of a highly-distorted mesh appear in linear elasticity both in the largest stable timestep that can be taken[18] and the accuracy of the solutions. Having already determined that $|e| = 1$ can be used in the SIPG formulation, we use it for all of the results that follow in this section on linear elasticity. Additionally, as solid mechanics problems often use low-order elements (nodal P1), we have restricted ourselves to using a 1^{st} order basis in order to have parity of polynomial basis degree. For all results shown, $\lambda = 1$ and $\mu = 1$.

Table 5-7. Stable timestep comparison

Mass matrix type	REBAR			DG	CG
	pen=100	pen=250	pen=500	pen=100	
Consistent (step size)	1.5e-3	1.5e-3	6.9e-4	2.4e-5	8.4e-4
Speedup	1.79x	1.79x	0.82x	0.03x	1.0x (ref.)
Lumped (step size)	1.2e-2	1.2e-2	5.3e-3	4.2e-5	1.4e-3
Speedup	**8.38x**	8.38x	3.74x	0.03x	1.0x (ref.)

Using 1^{st} order polynomial basis in GMLS.

Timestep size Table 5-7 shows a 8.38x stable timestep enlargement for our approach as compared to traditional CG-FEM for the extremely distorted annulus when using a lumped matrix. Timestep sizes are given for both the consistent and lumped form of the mass matrix and it can be noted that the timestep enlargement is less significant when using the consistent mass matrix. Additionally, there is an inverse relationship between the penalty parameter σ_e and the resulting stable timestep size.

Accuracy Next, we present the comparisons of the L^2 norm and H^1 seminorm errors calculated using the exact solution:

$$\vec{u} = \begin{bmatrix} \sin(x)\sin(y) \\ -\sin(x)\sin(y) \end{bmatrix}.$$

One thing that can be immediately noted from Figure 5-5 is that the error from shear level 2 (i.e. tf=1.795) in both norms for the GMLS+SIPG is smaller than the error from shear level 1 for CG-FEM. In fact, the error is 5.5x lower for the L^2 norm and 5.73x lower for the H^1 seminorm for the shear level 2 mesh using our approach. It can also be seen in Figure 5-5 that on each annulus, regardless of level of shear, optimal convergence rates are observed in each norm (Cf. the black dotted line).

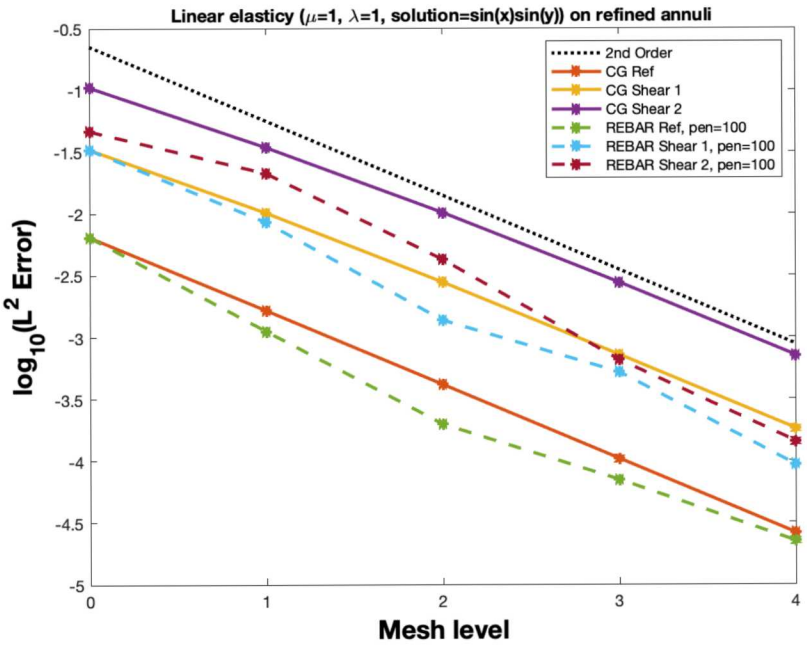

(a) L^2 error comparison

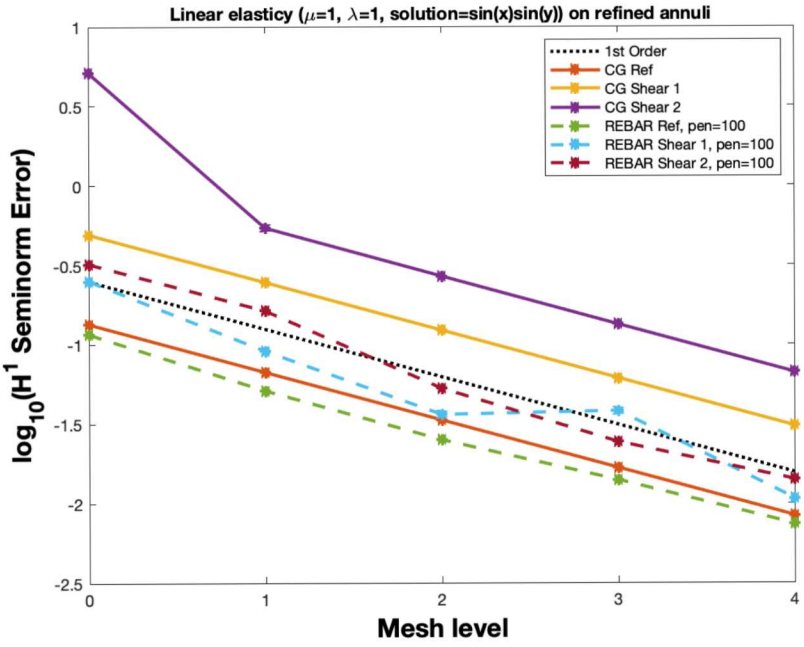

(b) H^1 seminorm error comparison

Figure 5-5. Comparison of errors for GMLS+SIPG vs. CG-FEM

5.4. Stokes Flow / incompressible

We now turn our attention towards the application of GMLS basis functions + SIPG for Stokes flow, which has an analog in linear elasticity in the incompressible regime. The manufactured solution is given by

$$\vec{u} = \begin{bmatrix} \cos(x)\sin(y) \\ -\sin(x)\cos(y) \end{bmatrix}, p = \sin(x)\sin(y)$$

and the viscosity is one.

Inf-sup stability It has already been shown numerically that our method converges optimally for elliptic problems. The significant difference between the previously discussed applications and Stokes flow is whether the formulation is inf-sup stable.

In Figure 5-6, a mesh is shown on which computed pressure solutions will be plotted. This mesh is relatively coarse, which is useful in visually identifying the characteristic checkerboard pattern often seen when inf-sup conditions are not satisfied by a formulation. We first consider equal order pairs of velocity and pressure bases. In CG-FEM this is a traditionally unstable pair (stabilization of pressure modes can be added), and we have no reason to expect inf-sup stability using these equal order pairs for the GMLS basis. As seen in Figure 5-7a, checkerboarding is indeed evident and can be handled gracefully in the SIPG framework by penalizing the jump in the pressure as in [7]. To see that this resolves the checkerboarding issue, cf. Figure 5-7b.

Next, we look at a 2^{nd} order velocity basis with a 1^{st} order pressure basis on the same coarse mesh, Figure 5-8a. While there certainly is no checkerboarding commensurate with what was seen in Figure 5-7a, it is not immediately obvious the fluctuations in the pressure field are in fact PDE solution error of a stable formulation and not a subtle checkerboarding phenomena. We refine the mesh twice and show the pressure field in Figure 5-8b. With refinement, there is less error in the approximation and it becomes immediately evident that the 2^{nd} order velocity basis with a 1^{st} order pressure basis is not displaying checkerboarding and therefore is likely (numerical demonstration, not a proof) inf-sup stable as a pair, similar to CG-FEM P2/P1 pairs.

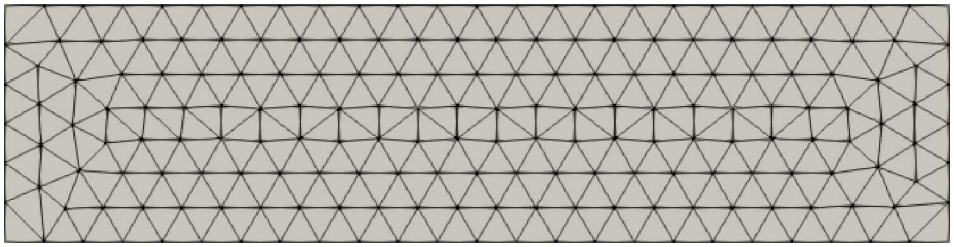

Figure 5-6. Triangular mesh of rectangular domain

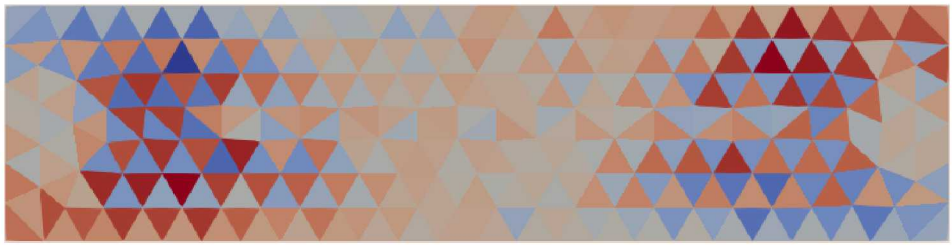

(a) Without additional stabilization
(checkerboarding, inf-sup instability)

(b) With pressure-jump stabilization [7]
(no checkerboarding)

Figure 5-7. Pressure plot using 1^{st} order basis for velocity, 1^{st} order for basis for pressure

Convergence To study the convergence rates of the velocity and pressure field using the GMLS+SIPG approach on Stokes flow, we consider two velocity/pressure pairs, namely $P1/P1$ and $P2/P1$. To confirm optimal convergence rates we use the method of manufactured solutions and plot the L^2 norm and H^1 seminorm errors for velocity and the L^2 norm for pressure, calculated on the rectangular domain in Figure 5-6.

Using a 1^{st} order basis for velocity and pressure, we see that it does even better than the expected 1^{st} order convergence in L^2 for velocity, L^2 for pressure, and H^1 seminorm for velocity, Cf. Figure 5-9. It isn't surprising that 1^{st} order may be surpassed and even 2^{nd} order achieved, given the alternate theoretical convergence rates given in [7] if additional criteria are satisfied. This is a demonstration that the SIPG stabilization of the equal order pair is consistent.

Using a 2^{nd} order basis for velocity and 1^{st} order basis for pressure, we see the expected 3^{rd} order convergence in L^2 for velocity and 2^{nd} order convergence in L^2 for pressure and H^1 seminorm for velocity 5-10. This adds support to the claim that the pairing of 2^{nd} order velocity with 1^{st} order pressure is inf-sup stable, as it is for CG-FEM.

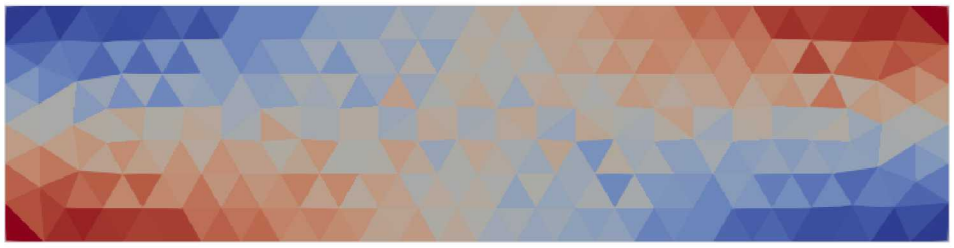

(a) Oscillations due to coarse mesh solution

(b) No oscillations with refinement

Figure 5-8. Pressure plot fine mesh using 2^{nd} order basis for velocity, 1^{st} order for basis for pressure without additional stabilization

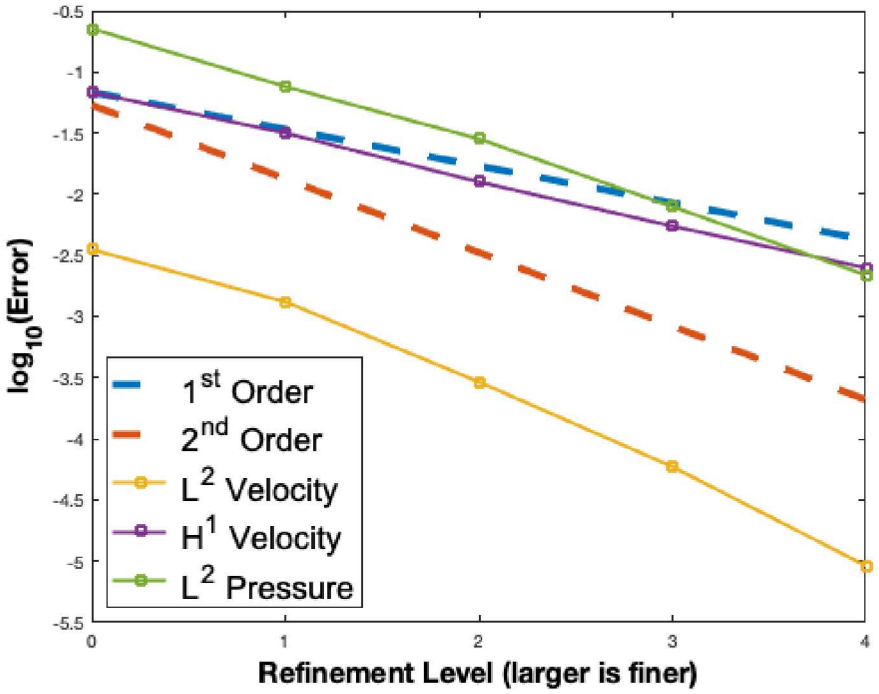

Figure 5-9. Plot of various norms for 1^{st} order basis for velocity and 1^{st} order basis for pressure (theoretical rate is 1 for all norms shown [7])

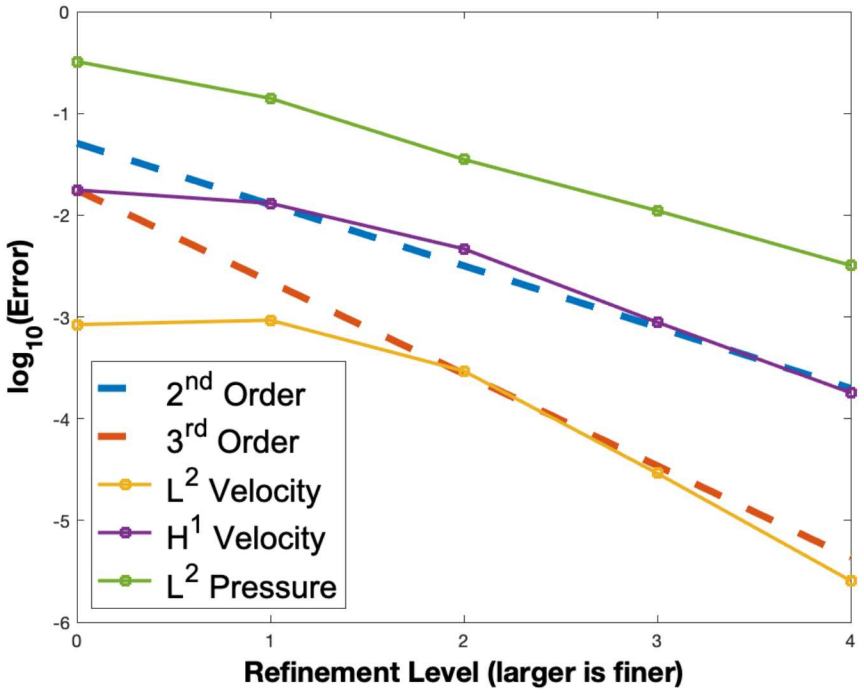

Figure 5-10. Plot of various norms for 2^{nd} order basis for velocity and 1^{st} order basis for pressure

6. CONCLUSIONS

We formulated a non-conforming finite element method which uses a new type of shape functions based on meshfree techniques. These shape functions are defined independently of the underlying mesh and break the dependence of the discretization quality on the mesh quality. Because the shape functions are piecewise polynomials, they can be easily incorporated into any standard non-conforming finite element framework.

We demonstrated numerically that the approach is stable and consistent, providing optimal asymptotic rates of convergence for the PDEs in all cases studied. Numerical results indicate that on a good quality mesh, the approach does not do as well as CG-FEM in an absolute error sense, although for poor quality meshes it is able to improve on the CG-FEM solution and deteriorates far less quickly with mesh quality.

This is a positive result given that nodal finite elements are a long-developed and well understood solution technique, whereas the approach we are developing is novel and there are a number of ways in which it could be incrementally improved (stabilization technique, GMLS kernel, search size, particle set choice, etc...).

In addition to an increased robustness of the resulting finite element method on ill-conditioned grids, we observe significant additional benefits including more relaxed definitions of the stabilization parameters as compared to traditional SIPG, resulting in an increased stable timestep size. While it is likely that $|e| = 1$ will not produce a stable result in the limit as $|e| \to 0$, numerical results indicate that the reliance of the stability terms on mesh size is significantly reduced.

Future work will consider extension of the approach to transient hyperelasticity and plasticity, using alternative stabilization techniques including variational multiscale (VMS).

7. ONGOING COLLABORATIONS

This project has an academic alliance with Prof. Arif Masud at the University of Illinois Urbana-Champaign and his graduate student, Marcelino Anguiano. During the course of this project, we have shared code with UIUC and had both members of the group visit Sandia to discuss technical details relating to applying VMS to the approach detailed in this report. The academic alliance will continue after the conclusion of this project, and the team at UIUC will continue to investigate and report on properties of the VMS approach using GMLS shape functions.

Quang Ha is a graduate student at Boston University who works under his advisor Paul Barbone. Quang visited Sandia for a considerable portion of the year that this project was taking place and is familiar with the internal workings of the Compadre Toolkit. As a result, he and his advisor will continue investigating properties of the SIPG approach using GMLS shape functions that were described in this report.

REFERENCES

[1] D. Arnold. An interior penalty finite element method with discontinuous elements. *SIAM Journal on Numerical Analysis*, 19(4):742–760, 1982.

[2] D. N. Arnold, F. Brezzi, B. Cockburn, and L. D. Marini. Unified analysis of discontinuous Galerkin methods for elliptic problems. 39:1749–1779, 2002.

[3] S. N. Atluri and T. Zhu. A new meshless local petrov-galerkin (MLPG) approach in computational mechanics. *Computational Mechanics*, 22(2):117–127, 1998.

[4] I. Babuška and A. Aziz. On the angle condition in the finite element method. *SIAM Journal on Numerical Analysis*, 13(2):214–226, 1976.

[5] T. Belytschko, Y. Y. Lu, and L. Gu. Element-free galerkin methods. *International Journal for Numerical Methods in Engineering*, 37(2):229–256, 1994.

[6] Pavel Bochev, Nathaniel A. Trask, Paul Kuberry, and Mauro Perego. Mesh-hardened finite element analysis through a generalized moving least-squares approximation of variational problems. In Ivan Lirkov and Svetozar Margenov, editors, *Large-Scale Scientific Computing. LSSC 2019.*, number 11958 in Lecture Notes in Computer Science, pages 67–75, Berlin, Heidelberg, 2020. Springer Verlag.

[7] Erik Burman and Peter Hansbo. A unified stabilized method for stokes and darcy's equations. *Journal of Computational and Applied Mathematics*, 198(1):35–51, 2007.

[8] Jiun-Shyan Chen, Michael Hillman, and Marcus Rüter. An arbitrary order variationally consistent integration for galerkin meshfree methods. *International Journal for Numerical Methods in Engineering*, 95(5):387–418, 2013.

[9] Jiun-Shyan Chen, Cheng-Tang Wu, Sangpil Yoon, and Yang You. A stabilized conforming nodal integration for galerkin mesh-free methods. *International Journal for Numerical Methods in Engineering*, 50(2):435–466, 2001.

[10] P. Ciarlet. *The Finite Element Method for Elliptic Problems*. SIAM Classics in Applied Mathematics. SIAM, Philadelphia, 2002.

[11] B. Cockburn, B. Dong, and J. Guzmán. Optimal convergence of the original dg method for the transport-reaction equation on special meshes. *SIAM Journal on Numerical Analysis*, 46(3):1250–1265, 2008.

[12] H. Carter Edwards, Christian R. Trott, and Daniel Sunderland. Kokkos: Enabling manycore performance portability through polymorphic memory access patterns. *Journal of Parallel and Distributed Computing*, 74(12):3202 – 3216, 2014. Domain-Specific Languages and High-Level Frameworks for High-Performance Computing.

[13] Yekaterina Epshteyn and Béatrice Rivière. Estimation of penalty parameters for symmetric interior penalty galerkin methods. *Journal of Computational and Applied Mathematics*, 206(2):843 – 872, 2007.

[14] M.F. Harwick, R.L. Clay, P.T. Boggs, E.J. Walsh, A.R. Larzelere, and A. Altshuler. Dart system analysis. Technical Report SAND2005-4647, Sandia National Laboratories, 2005.

[15] Paul A. Kuberry, Peter A. Bosler, and Nathaniel A. Trask. Compadre toolkit. [Computer Software] https://doi.org/10.11578/dc.20190411.1, jan 2019.

[16] Davoud Mirzaei, Robert Schaback, and Mehdi Dehghan. On generalized moving least squares and diffuse derivatives. *IMA Journal of Numerical Analysis*, 32(3):983–1000, 2012.

[17] B. Nayroles, G. Touzot, and P. Villon. Generalizing the finite element method: Diffuse approximation and diffuse elements. *Computational Mechanics*, 10(5):307–318, 1992.

[18] Nathan M Newmark. A method of computation for structural dynamics. *Journal of the engineering mechanics division*, 85(3):67–94, 1959.

[19] M. A. Puso, J. S. Chen, E. Zywicz, and W. Elmer. Meshfree and finite element nodal integration methods. *International Journal for Numerical Methods in Engineering*, 74(3):416–446, 2008.

[20] Jonathan Richard Shewchuk. What is a good linear element? interpolation, conditioning, and quality measures. In *In 11th International Meshing Roundtable*, pages 115–126, 2002.

[21] Nathaniel Trask, Mauro Perego, and Pavel Bochev. A high-order staggered meshless method for elliptic problems. *SIAM Journal on Scientific Computing*, 39(2):A479–A502, 2017.

[22] The Trilinos Project Team. *The Trilinos Project Website*.

[23] Holger Wendland. *Scattered data approximation*, volume 17. Cambridge university press, 2004.

[24] Mary Fanett Wheeler. An elliptic collocation-finite element method with interior penalties. *SIAM Journal on Numerical Analysis*, 15(1):152–161, 1978.

DISTRIBUTION

[redacted]

Name	Org.	Sandia Email Address
Pavel Bochev	1400	pbboche@sandia.gov
Scott Collis	1400	sscoll@sandia.gov
James Stewart	1440	jrstewa@sandia.gov
Paul Kuberry	1442	pakuber@sandia.gov
Michael Parks	1442	mlparks@sandia.gov
Nathaniel Trask	1442	natrask@sandia.gov
Corey Cruz	1500	cacruz@sandia.gov
Martin Heinstein	1500	mwheins@sandia.gov
Walter Witkowski	1540	wrwitko@sandia.gov
Kendal Pierson	1542	khpiers@sandia.gov
Scott Klenke	1550	seklenk@sandia.gov
Aaron Brundage	1554	albrund@sandia.gov
Jacob Koester	1554	jkoeste@sandia.gov
Joe Bishop	1556	jebisho@sandia.gov
Joel Lash	4000	jslash@sandia.gov
Technical Library	01177	libref@sandia.gov

Sandia National Laboratories is a multimission laboratory managed and operated by National Technology & Engineering Solutions of Sandia LLC, a wholly owned subsidiary of Honeywell International Inc., for the U.S. Department of Energy's National Nuclear Security Administration under contract DE-NA0003525.