

ADAPTIVE MESH REFINEMENT FOR OBSTACLE PROBLEMS*

G. STEFANO FOCESATTO[†] AND ED BUELER[‡]

Abstract. Free-boundary problems posed as variational inequalities, including obstacle problems, appear in many scientific and engineering applications. In their finite element (FE) solution, localization of the free boundary may be a primary goal, and the geometrical localization error often dominates the overall numerical error. In this paper we implement, using the Firedrake FE library, new parallel adaptive mesh refinement strategies which generate accurate, high-resolution free boundaries through h -refinement. We evaluate three approaches: (i) a tag-and-refine unstructured dilation operator method using discrete adjacency to the free boundary, (ii) a tag-and-refine method based on variable-coefficient diffusion, which thresholds a diffused active-set indicator function, and (iii) a metric-based mesh adaptation method which averages an anisotropic, Hessian-derived Riemannian metric with an isotropic metric computed from the diffused indicator in (ii). For (i) and (ii) classical *a posteriori* error estimators must be added within the computed inactive sets to attain convergence. These methods are evaluated, versus mesh complexity, by norm error and by geometrical localization using Jaccard distances for active sets. Applications include classical Laplacian obstacle problems and a shallow ice flow problem for predicting glaciation.

Key words. variational inequality, free boundary, obstacle problem, finite element, adaptive mesh refinement, glaciers

MSC codes. 35J85, 35R35, 65N50

1. Introduction. The classical Laplacian obstacle problem [25] finds the equilibrium vertical displacement u of an elastic membrane over some domain $\Omega \subset \mathbb{R}^2$. The membrane, attached with displacement g at the fixed boundary $\partial\Omega$, is subjected to an applied force f , but it is also constrained to be above a given obstacle ψ . The strong formulation is thus a complementarity problem over Ω :

$$(1.1a) \quad -\nabla^2 u - f \geq 0$$

$$(1.1b) \quad u - \psi \geq 0$$

$$(1.1c) \quad (-\nabla^2 u - f)(u - \psi) = 0$$

From a solution of (1.1), or rather of its weak form (below), we may identify the inactive and active sets, and the free boundary:

$$(1.2) \quad I_u = \{x \in \Omega : u(x) > \psi(x)\}, \quad A_u = \Omega \setminus I_u, \quad \Gamma_u = \Omega \cap \partial I_u.$$




FIG. 1. Solution, as wireframe, to a problem with a (hemi)spherical obstacle.

*Submitted to the editors DATE (draft August 15, 2025).

[†]Dept. of Mathematics and Statistics, University of Alaska Fairbanks, USA

[‡]corresponding author: elbueler@alaska.edu

For the example shown in Figure 1, the obstacle ψ is an upper hemisphere, the active set A_u (white mesh) is a disc, and the free boundary Γ_u is a circle. Note that u solves the Poisson equation $-\nabla^2 u = f$ on I_u ; this “interior condition” of the problem is the black mesh in the Figure. Note that both Dirichlet ($u = \psi$) and Neumann ($\partial u / \partial n = \partial \psi / \partial n$) conditions apply along the unknown free boundary Γ_u .

Another physical interpretation of problem (1.1) is that u models the water pressure, which cannot go below zero ($\psi = 0$), in a porous dam [2, for example]; see Example 4.3. Section 6 will present a different obstacle problem application with a highly-nonlinear operator. There the solution is the surface elevation of a glacier, which is constrained to be above the bedrock elevation on which the glacier sits.

Problem (1.1) has a weak formulation which is a variational inequality (VI) [25] over a Sobolev space. Let $\Omega \subset \mathbb{R}^d$ be the domain, $d \geq 1$, let $\mathcal{X} = H^1(\Omega)$ [16], and suppose $\psi \in \mathcal{X} \cap C(\bar{\Omega})$. Let $g : \partial\Omega \rightarrow \mathbb{R}$ be continuous, with $g \geq \psi|_{\partial\Omega}$, and define

$$(1.3) \quad \mathcal{K} = \{u \in \mathcal{X} : u \geq \psi \text{ and } u|_{\partial\Omega} = g\}$$

as the admissible subset, which is closed and convex in $\mathcal{X}$. For $f \in L^2(\Omega)$, the VI formulation of (1.1) finds $u \in \mathcal{K}$ so that

$$(1.4) \quad \int_{\Omega} \nabla u \cdot \nabla (v - u) \geq \int_{\Omega} f(v - u) \quad \text{for all } v \in \mathcal{K}.$$

In Section 2 we will recall the theory of such VIs, and extend the theory of their finite element (FE) approximation. We will generalize (1.4) from elliptic bilinear forms like (1.4) to coercive nonlinear operators over Banach spaces. In an FE approximation of such a VI the numerical solution u_h will solve the same weak form, but over a finite-dimensional admissible set constructed on a mesh $\mathcal{T}_h$. Similarly to Cea’s lemma for PDEs [15], the norm errors $\|u - u_h\|$ can be bounded *a priori*, which we do by extending the Falk [17] technique to nonlinear operators. This will show how norm errors are controlled by FE space approximation properties, as usual, but subject to admissibility concerns, and with separation of active-set and inactive-set errors.

Adaptive mesh refinement (AMR) uses *a posteriori* information from the numerical solution to strategically add mesh elements to increase the resolution and reduce the numerical error. However, for VI problems the simulation goal is often the accurate approximation of the sets in (1.2), attainability of which is problem-dependent.

Example 1.1. Suppose $\Omega = (-1, 1)$, $\psi(x) = 1 - x^2$, and $g(x) = 0$, but let f be constant: $f(x) = \alpha$. The exact solution u of (1.4) is now easily calculated: $u(x) = \psi(x)$ if $\alpha \leq 2$ and $u(x) = 0.5\alpha\psi(x) > \psi(x)$ if $\alpha > 2$. Thus for $\alpha \geq 2$ we have $A_u = \Omega$ and $I_u = \emptyset$, while if $\alpha < 2$ then $A_u = \emptyset$ and $I_u = \Omega$. This example shows that the sets (1.2) are not continuous functions of the data f of the problem. An easy modification of this example shows that the sets are also not continuous functions of ψ .

The $\alpha = 2$ case of Example 1.1 is *degenerate* [25], that is, the unconstrained solution happens to match the obstacle on a nonempty open set. Usable convergence results for FE approximations of the solution-dependent sets A_u, I_u, Γ_u will usually depend upon problem nondegeneracy, and our examples are accordingly non-degenerate. Furthermore, the effectiveness of different AMR strategies for non-degenerate VI problems depends strongly upon the measure (area or volume) of the active and inactive sets. This observation largely motivates the *a posteriori* approaches of this paper. For example, in certain problems, those elements which are significantly interior to the active set require no further computation or refinement. For such problems, with

examples given in Sections 5 and 6, if a solution is desired at higher resolution within the active set then this can be computed in post-processing by arbitrary interpolation of the obstacle data ψ .

Example 1.2. Examples in Section 5 include three classical obstacle problems over square domains, and Figure 2 shows their active sets in black. For the left-hand problem, with a small active set and a long free boundary, our AMR methods refine a large fraction of the elements, namely the many elements which are close to the free boundary. The middle problem has a known exact solution; Example 5.3 shows convergence rates. For the right-hand problem, a clear performance benefit of our techniques, relative to uniform refinement, comes from avoiding refinement in the active set, an efficiency also exploited by the glaciation application in Section 6.




FIG. 2. The area (measure) of the active set (black) can vary from small to large (left to right); the middle image matches Figure 1.

In this work we consider only P_1 element spaces over meshes of triangles or tetrahedra, and only h -refinement is addressed. However, VIs can be solved using p -refinement and higher-order elements, once nontrivial penalty-type modifications are made to the VI [24], but this is not attempted here. Also, while the classical obstacle problem (1.4) is equivalent to constrained minimization of a scalar objective, our analysis of FE errors for VI problems will *not* require such an objective, and our refinement strategies do not exploit one if available.

Three AMR methods for VIs are introduced and detailed in Section 4. Our implementations use the Firedrake FE library [31] and generate conforming meshes with no hanging nodes. The first two methods are of tag-and-refine type, only differing by which elements are tagged, with skeleton-based refinement (SBR) [29] applied after tagging. These two methods also require complementary refinement of the PDE problem in the inactive set to achieve convergence; see Section 3. The third method uses the Netgen and Animate [33] libraries for goal-oriented, metric-based mesh adaptation. Here is a high-level view of the new methods:

UDO: The unstructured dilation operator method discretely identifies elements adjacent to the computed free boundary, employing a graph-based approach to tag neighboring elements for refinement. It generalizes the image processing operation of dilation [30] to unstructured meshes.

VCD: The variable-coefficient diffusion method starts from a node-wise indicator function for the current computed active set. This indicator becomes the initial iterate in a single step of a time-dependent heat equation problem, which

smooths the indicator about the free boundary. This smoothed indicator is then thresholded for element tagging and refinement.

AVM: The averaged-metric method computes an intermediate representation of the size, shape, and orientation of a new mesh, namely as a tensor-valued Riemannian metric [3]. Here the metric is an average of an anisotropic metric, from the Hessian of the computed solution [33], and an isotropic metric derived from the diffused active set indicator of the above VCD method. This method can maintain mesh complexity as it simultaneously resolves the free boundary and reduces errors in the inactive set.

Note that the UDO and VCD methods generally produce similar results, but their motivation, and their control parameters, are sufficiently different to justify separate presentation.

Example meshes are shown in Figure 3, generated by three levels of refinement using the above three methods, starting from a coarse uniform mesh, on the obstacle problem shown in Figure 2 (middle). All three schemes quickly concentrate effort around an accurately-localized free boundary, augmented by refinement as needed in the inactive set (Section 3). This kind of AMR accelerates convergence and reduces unnecessary computation.




FIG. 3. Meshes from UDO (left), VCD (middle), and AVM (right) methods.

AMR for VI problems has only been lightly explored in the literature. The first published analysis may be [2], giving an error bound for the classical obstacle problem in terms of local functionals associated with each element. The monograph by Suttmeier [32] covers a broader class of problems, including elasticity. For the classical obstacle problem, the constructable error estimators in these works require heuristic assumptions which may not hold in general. (See inequality (42) in [2], and the approximation “ $(u - \psi)\lambda_h \approx 0$ ” in [32].) To our knowledge these approaches are not found in publicly-available implementations, nor are they as efficient as our strategies for computing high-resolution approximations to free boundaries.

Our focus in this paper is on AMR performance, not solver performance. For all examples we used a fixed, VI-adapted, reduced-space Newton method with line search [7], implemented in PETSc [5]. Note that since the constraint $u \geq \psi$ makes VI problems nonlinear, an iterative solver is required even if the operator is linear as in (1.4). Such a numerical method cannot converge quadratically until the active and inactive sets stabilize on the given mesh. Then convergence will occur in one additional iteration, for a linear operator, or otherwise in a few iterations for a well-behaved nonlinear operator.

In summary, here are two principles for the AMR methods of this paper:

1. Relative to uniform refinement, they exhibit significant improvements in convergence rate. This is measured by norms, or especially by free-boundary localization (geometrical) error (Section 5), per mesh degree of freedom.
2. Their implementations (github.com/StefanoFochesatto/viamr) within the Firedrake [26] FE library are parallel, well-documented, and easy-to-use.

The paper is organized as follows: Section 2 provides *a priori* norm bounds for FE methods applied to VI problems; aspects of this material are new. Section 3 addresses *a posteriori* error estimators which can be applied in the inactive set. Section 4 describes the three new AMR methods in more detail. Sections 5 and 6 compare and discuss their performance on classical obstacle model problems and in a realistic glacier application. Table 1 states the few abbreviations used herein.

AMR	adaptive mesh refinement	GR	gradient recovery
AVM*	averaged-metric	PDE	partial differential equation
BR	Babuška–Rheinboldt	SBR	skeleton-based refinement
CG	continuous Galerkin	UDO*	unstructured dilation operator
DG	discontinuous Galerkin	VCD*	variable-coefficient diffusion
FE	finite element	VI	variational inequality

TABLE 1

Abbreviations used in this paper. Stars indicate the new AMR methods.

2. Variational inequalities and their finite element approximations. We consider unilateral obstacle problems in Banach spaces. Let $\Omega \subset \mathbb{R}^d$, $d \geq 1$, be a bounded, polygonal domain. Let $\mathcal{X} = W^{1,p}(\Omega)$, $p > 1$, be the Sobolev space of measurable functions with p th-integrable weak gradients [16]. We will assume continuous problem data, with well-defined point values, so suppose $\psi \in \mathcal{X} \cap C(\bar{\Omega})$ and $g \in C(\partial\Omega)$ satisfy $g \geq \psi|_{\partial\Omega}$. Define the closed and convex admissible subset

$$(2.1) \quad \mathcal{K} = \{v \in \mathcal{X} : v \geq \psi \text{ and } v|_{\partial\Omega} = g\} \subset \mathcal{X},$$

same as in (1.3). Observe that generally $\psi \notin \mathcal{K}$.

Let $\mathcal{X}'$ be the dual space of $\mathcal{X}$, and denote the application of $\omega \in \mathcal{X}'$ to $v \in \mathcal{X}$ by $\omega[v] \in \mathbb{R}$. The norm on $\mathcal{X}$ is denoted $\|\cdot\|$, and for $\mathcal{X}'$ the norm is $\|\omega\|' = \sup_{\|v\|=1} |\omega[v]|$. Let $F : \mathcal{K} \rightarrow \mathcal{X}'$ be a given operator, generally nonlinear, and let $\ell \in \mathcal{X}'$ be given. (While F in Example 2.1 is defined on all of $\mathcal{X}$, the problem in Section 6 shows how F might be defined only on $\mathcal{K}$.) The VI associated to this data, a unilateral obstacle problem, finds $u \in \mathcal{K}$ so that

$$(2.2) \quad F(u)[v - u] \geq \ell[v - u] \quad \text{for all } v \in \mathcal{K}.$$

The problems in this paper can be analyzed within the framework of coercivity and Lipschitz continuity. We say F is q -coercive, $1 < q < \infty$, if there is $\alpha > 0$ so that

$$(2.3) \quad (F(v) - F(w))[v - w] \geq \alpha \|v - w\|^q$$

for all $v, w \in \mathcal{K}$. Note that if F is q -coercive then it is also strictly monotone: $(F(v) - F(w))[v - w] > 0$ for $v \neq w$. Let $B_R(0)$ be the open ball at $0 \in \mathcal{X}$ of radius $R > 0$. We say F is Lipschitz on bounded subsets if there is $C(R) > 0$ so that

$$(2.4) \quad \|F(v) - F(w)\|' \leq C(R) \|v - w\|$$

for all $v, w \in B_R(0) \cap \mathcal{K}$. If F satisfies (2.4) then it is continuous. From coercivity, strict monotonicity, and continuity of F it follows that a unique solution to (2.2) exists [25, Corollary III.1.8].

Example 2.1. In the classical obstacle problem (1.4), over $\mathcal{X} = H^1(\Omega) = W^{1,2}(\Omega)$, $F(u)[v] = \int_{\Omega} \nabla u \cdot \nabla v \, dx$. This bilinear operator is 2-coercive because the Laplacian is uniformly elliptic [16], and it is Lipschitz over $\mathcal{X}$ with constant $C = 1$.

Example 2.2. In the glacier model of Section 6, when the bedrock is flat and the surface mass balance is independent of elevation the VI problem (6.1) uses a 4-Laplacian operator over transformed thicknesses $u \in \mathcal{X} = W^{1,4}(\Omega)$: $F(u)[v] = \int_{\Omega} \Gamma |\nabla u|^2 \nabla u \cdot \nabla v \, dx$ where $\Gamma > 0$ is constant. This operator is 4-coercive [22, for example], and Lipschitz on bounded subsets of $\mathcal{X}$.

Note that if the inequality constraint in (2.2) were absent then the residual of the solution u would be zero ($F(u) - \ell = 0$); this is the PDE case. However, for VI (2.2) we only have that $F(u) - \ell = 0$ a.e. within an unknown inactive set I_u . The residual $F(u) - \ell \in \mathcal{X}'$ might be highly-irregular in the active set A_u , but it is nonnegative. In fact, the following lemma states the weak complementarity property associated to obstacle problems (2.2); compare strong-form complementarity (1.1).

LEMMA 2.3. [25, Theorem II.6.9]. Suppose $u \in \mathcal{K}$ solves (2.2). Then $F(u) - \ell = d\mu_u$ is a positive Radon measure supported in A_u . Thus for $w \in \mathcal{X}$ we have

$$(2.5) \quad (F(u) - \ell)[w] = \int_{A_u} w \, d\mu_u.$$

Now let $\mathcal{T}_h$ be a shape-regular mesh partition (triangulation, etc.) of Ω [1, 15]. Let $\mathcal{X}_h \subset \mathcal{X} \cap C(\bar{\Omega})$ be a conforming finite-dimensional FE subspace over $\mathcal{T}_h$. (Our examples will be piecewise-linear over triangles and tetrahedra: $\mathcal{X}_h = P_1$.) Assume that there is $g_h \in \mathcal{X}_h$ such that $g_h = g$ along $\partial\Omega$. Let $\psi_h \in \mathcal{X}_h$ be the FE obstacle, which satisfies the compatibility requirement $\psi_h \leq g_h$ along $\partial\Omega$. Define the (nonempty) FE admissible set

$$(2.6) \quad \mathcal{K}_h = \{v_h \in \mathcal{X}_h : v_h \geq \psi_h \text{ and } v_h|_{\partial\Omega} = g_h|_{\partial\Omega}\}.$$

Our FE method seeks $u_h \in \mathcal{K}_h$ satisfying a VI problem which approximates (2.2):

$$(2.7) \quad F(u_h)[v_h - u_h] \geq \ell[v_h - u_h] \quad \text{for all } v_h \in \mathcal{K}_h.$$

The same argument given for (2.2) shows that this has a unique solution u_h . Define

$$(2.8) \quad I_u^h = \{x \in \Omega : u_h(x) > \psi_h(x)\}, \quad A_u^h = \Omega \setminus I_u^h, \quad \Gamma_u^h = \Omega \cap \partial I_u^h,$$

the numerical sets corresponding to (1.2), defined *a posteriori* from solving (2.7).

Because $\mathcal{K}_h \subset \mathcal{X}_h \subset \mathcal{X}$, we might regard (2.7) as a conforming FE method for (2.2). However, this is subtle for an obstacle problem as it depends on the relationship between ψ and ψ_h . If ψ_h is an interpolant of ψ then $\mathcal{K}_h \approx \mathcal{K}$ in some sense. Actually, one should distinguish three levels of increasingly “conforming” admissibility:

- i) $\mathcal{K}_h \not\subset \mathcal{K}$
- ii) $\mathcal{K}_h \subset \mathcal{K}$
- iii) $\mathcal{K}_h = \mathcal{K} \cap \mathcal{X}_h$

Ciarlet [14, Figure 5.1.3] observed early on that situation i) generally applies for an interpolated obstacle $\psi_h = \Pi_h \psi$, e.g. for $\mathcal{X}_h = P_1$ and when ψ is not convex. Situation

ii) holds when $\psi_h \geq \psi$, which can be imposed by using a monotone injection operator, e.g. “ $\psi_h = R^\oplus \psi$ ” in the notation from [13]. (For related ideas, see [20] and the proof of Theorem 5.1.2 in [14].) The strongest condition iii) holds if $\psi_h = \psi$ exactly, for example when $\psi = 0$ in the porous dam problems considered by [2], and in Section 6.

In any case, the obstacle ψ is given *as data* for VI problem (2.2) over $\mathcal{K}$. As illustrated in Section 5, point values of ψ can be evaluated as desired to better-represent u_h within the numerical active set A_u^h . Over elements where there is active-set correctness, namely $K \in \mathcal{T}_h$ such that $K \subset A_u \cap A_u^h$, the error $e = u - u_h = \psi - \psi_h$ can simply be regarded as unimportant. Alternatively, the approximation $\psi_h \approx \psi$ can be improved as needed by better interpolation of the data ψ . As the goal for AMR is to more-accurately solve the FE problem, refinement within a stabilized active set, once an accurate free boundary has been found, is wasted effort.

These ideas are already implicit in *a priori* bounds for FE error in VI problems. The following theorem generalizes the Falk bound [17]; see also [14, Theorem 5.1.1]. It can be extended further to address $F_h \approx F$ [12, Theorem 6.3], but we will not need such an operator approximation for our examples.

THEOREM 2.4. *For $1 < q < \infty$, define the conjugate exponent $q' = q/(q-1)$. Assume that F is q -coercive and Lipschitz on bounded subsets of its domain. Suppose $u \in \mathcal{K}$ solves (2.2) and $u_h \in \mathcal{K}_h$ solves (2.7). Let $R_h = \max\{\|u\|, \|u_h\|\}$. Then there is a constant $c(R_h) > 0$, not otherwise depending on u or u_h , so that*

$$(2.9) \quad \|u - u_h\|^q \leq \frac{2}{\alpha} \left(\inf_{v \in \mathcal{K}} \int_{A_u} (v - u_h) d\mu_u + \inf_{v_h \in \mathcal{K}_h} \int_{A_u} (v_h - \psi) d\mu_u + c(R_h) \inf_{v_h \in \mathcal{K}_h} \|v_h - u\|^{q'} \right).$$

Proof. For arbitrary $v \in \mathcal{K}$ and $v_h \in \mathcal{K}_h$, rewrite (2.2) and (2.7) as $F(u)[u] \leq F(u)[v] + \ell[u - v]$ and $F(u_h)[u_h] \leq F(u_h)[v_h] + \ell[u_h - v_h]$, respectively. It follows from these inequalities, and q -coercivity of F , that

$$(2.10) \quad \begin{aligned} \alpha \|u - u_h\|^q &\leq (F(u) - F(u_h)) [u - u_h] \\ &= F(u)[u] + F(u_h)[u_h] - F(u)[u_h] - F(u_h)[u] \\ &\leq F(u)[v] + \ell[u - v] + F(u_h)[v_h] + \ell[u_h - v_h] \\ &\quad - F(u)[u_h] - F(u_h)[u] \\ &= F(u)[v - u_h] - \ell[v - u_h] + F(u_h)[v_h - u] - \ell[v_h - u] \\ &= (F(u) - \ell) [v - u_h] + (F(u) - \ell) [v_h - u] \\ &\quad + (F(u) - F(u_h)) [u - v_h] \end{aligned}$$

Since $u, u_h \in B_{R_h} = \{w \in \mathcal{X} : \|w\| \leq R_h\}$, by the Lipschitz assumption (2.4) there is $C(R_h) > 0$ so that the last term from (2.10) has bound

$$(2.11) \quad (F(u) - F(u_h)) [u - v_h] \leq C(R_h) \|u - u_h\| \|u - v_h\|.$$

Now use Young’s inequality with $\epsilon > 0$ [16, Appendix B.2] on (2.11). We have:

$$(2.12) \quad \begin{aligned} \alpha \|u - u_h\|^q &\leq (F(u) - \ell) [v - u_h] + (F(u) - \ell) [v_h - u] \\ &\quad + C(R_h) \left(\epsilon \|u - u_h\|^q + \tilde{C}(\epsilon) \|u - v_h\|^{q'} \right) \end{aligned}$$

where $\tilde{C}(\epsilon) = (\epsilon q)^{-q'/q} q'^{-1}$. Choose $\epsilon > 0$ so that $C(R_h)\epsilon \leq \alpha/2$. Then

$$(2.13) \quad \frac{\alpha}{2} \|u - u_h\|^q \leq (F(u) - \ell) [v - u_h] + (F(u) - \ell) [v_h - u] + C(R_h)\tilde{C}(\epsilon) \|u - v_h\|^{q'}$$

Apply Lemma 2.3 and take infimums to show (2.9). $\square$

Consider the unconstrained PDE case of (2.9), namely when $A_u = \emptyset$. In this case the bound is simply Cea's lemma (quasi-optimality) for q -coercive operators: $\|u - u_h\| \lesssim \inf_{v_h \in \mathcal{K}_h} \|v_h - u\|^{1/(q-1)}$. It is standard in FE theory [1, 15] to address this bound by solution regularity and interpolation theory.

If $\psi_h \geq \psi$ then $\mathcal{K}_h \subset \mathcal{K}$, and so the first term on the right of (2.9) can be replaced by zero. In this case bound (2.9) adds a single term to Cea's lemma, which is nonzero when $v_h \in \mathcal{K}_h$ is blocked by ψ_h from descending close to the continuum obstacle ψ in the active set A_u . Thus this term is large if ψ_h is substantially above ψ in A_u , or if $d\mu_u$ is large in A_u . However, if the FE method has generated a numerical active set which is accurate, $A_u^h \approx A_u$, the errors in the active set are irrelevant because the data ψ is available to the solver.

A priori error bound (2.9) is further clarified in the $p = 2$ and $q = q' = 2$ case, the classical obstacle problem. The bound can then be split between integrals over the (exact) active set and inactive sets.

COROLLARY 2.5. *Suppose all the hypotheses of Theorem 2.4. Assume that $p = q = 2$, and that $\psi \in C^1(\Omega)$. Up to constants which depend on u and u_h , we may write the a priori bound with four integrals,*

$$(2.14) \quad \|u - u_h\|^2 \lesssim \inf_{v \in \mathcal{K}} \int_{A_u} (v - u_h) d\mu_u + \inf_{v_h \in \mathcal{K}_h} \int_{A_u} (v_h - \psi) d\mu_u \\ + \inf_{v_h \in \mathcal{K}_h} \left(\int_{A_u} |\nabla v_h - \nabla \psi|^2 dx + \int_{I_u} |\nabla v_h - \nabla u|^2 dx \right)$$

Proof. Apply Poincaré's inequality to the final term in (2.9). $\square$

Consider the first two integrals in bound (2.14) from the point of view of AMR. Their size is determined partly by the action of F on the continuum obstacle, and partly by the source term ℓ . However, because they are over the active set A_u , neither integral, nor the third A_u integral, requires mesh refinement so as to improve the quality of the FE solution, *as long as the mesh and solver have accurately located the free boundary*. This observation suggests why the primary goal of AMR for VIs should be to generate close approximation $\Gamma_u^h \approx \Gamma_u$. Computation expended on better approximation of the data ψ in A_u can be avoided if the free boundary is accurately resolved. This is easiest to exploit within the class of unilateral obstacle problems described in Appendix A, which includes the $\psi = 0$ case of the classical obstacle problem (1.4) and the glacier problem of Section 6. In such problems one may preprocess the data, noting all the areas where the source term is negative, and then systematically avoid unnecessary active set refinement.

However, in order for $\Gamma_u^h \approx \Gamma_u$ to be an accurate approximation, the interior condition of the VI, over I_u , must be accurately solved. (This gives the significance of the final integral in (2.14).) Our AMR approaches will systematically refine on both sides of the computed free boundary Γ_u^h , and generally in the inactive set I_u^h .

3. *A posteriori* error estimation in the inactive set. For $u \in \mathcal{X}$ solving VI (2.2), an interior condition PDE holds in the inactive set I_u [25]. Our AMR strategies

(Section 4) do *a posteriori* refinement in the geometric vicinity of the free boundary, but we will also need to exploit *a posteriori* error indicators for the interior condition to achieve convergence. We consider two such PDE-type indicators.

Example 3.1. Suppose $\mathcal{X}_h = \text{CG}_k$ is the continuous, piecewise-polynomial FE space of degree k , and consider $u_h \in \mathcal{X}_h$. Noting ∇u_h is discontinuous, but well-defined on each element, define $\mathcal{Y}_h = \text{DG}_{k-1}$ for the (scalar) discontinuous space with polynomial degree $k-1$, so that $\nabla u_h \in \mathcal{Y}_h^d$ (vector-valued). Suppose there is a linear operator $G : \mathcal{X}_h \rightarrow \mathcal{X}_h^d$, into the continuous space, called *gradient recovery* (GR) [1, Chapter 4], so that $G(u_h) \approx \nabla u_h$ in some sense. Over an element $K \in \mathcal{T}_h$, the corresponding error indicator $\eta_K \geq 0$ is then

$$(3.1) \quad \eta_K^2 = \int_K |G(u_h) - \nabla u_h|^2.$$

We also define $\eta^2 = \sum_{K \in \mathcal{T}_h} \eta_K^2$. For certain gradient recovery methods G , when they are applied to the Poisson equation, we find that $\eta \sim |u - u_h|_{H^1}$ (energy norm) as $h \rightarrow 0$ [1, Theorem 4.4].

Our application of GR in Section 6 will simply use orthogonal projection in L^2 for the map G . That is, $G(u_h) \in \mathcal{X}_h^d$ is defined to be the minimizer of

$$(3.2) \quad J(w) = \int_{\Omega} |w - \nabla u_h|^2 dx.$$

In this application $\mathcal{X}_h = \text{CG}_1$, ∇u_h is in vector-valued DG_0 , and $G(u_h)$ is in $\mathcal{X}_h^d$.

The second indicator, applied in Section 5 to classical problem (1.4), is a well-known *explicit error estimator* [1, Chapter 2].

Example 3.2. Suppose $u \in H^1(\Omega)$ solves the weak form Poisson equation $a(u, v) = \int_{\Omega} f v dx$ for all $v \in H_0^1(\Omega)$, with $a(u, v) = \int_{\Omega} \nabla u \cdot \nabla v dx$ and $u = g$ on $\partial\Omega$. Suppose that $u_h \in \mathcal{X}_h = \text{CG}_1$ solves the corresponding finite-dimensional weak form. For each element $K \in \mathcal{T}_h$, define n_K as the unit outward normal vector on ∂K . For a pair of elements L, K incident to an edge γ , and any vector field Z with traces Z_L, Z_K on γ , let $\llbracket Z \cdot n \rrbracket = Z_L \cdot n_L + Z_K \cdot n_K$ be the jump of Z on γ . Given the (strong) residual $r(u_h) = -\nabla^2 u_h - f$ of the Poisson equation, well-defined within each element K , the Babuška–Rheinboldt (BR) [4] error estimator is

$$(3.3) \quad \eta_K^2 = h_K^2 \int_K |r(u_h)|^2 dx + \frac{h_K}{2} \sum_{\gamma \in \partial K \setminus \partial\Omega} \int_{\gamma} \llbracket \nabla u_h \cdot n \rrbracket^2 dS,$$

where h_K is the diameter of K . It can be shown [1, Chapter 2] that the energy error is bounded by $\eta^2 = \sum_{K \in \mathcal{T}_h} \eta_K^2$:

$$(3.4) \quad |u - u_h|_{H^1}^2 = \int_{\Omega} |\nabla(u - u_h)|^2 dx \leq C\eta^2.$$

Similarly the L^2 error can be bounded by an estimator, which replaces the powers in (3.3) by h_K^4, h_K^3 , respectively [1, Section 2.4].

For the results in Sections 5 and 6, whether η_K is computed as in Example 3.1 or 3.2, we will treat the values η_k as local element-wise error estimators. The set $\{\eta_K\}$ will be thresholded to providing tagging of elements for refinement; see [6, Section 4.2]

for a discussion of techniques. In our applications, all (computed) inactive elements satisfying $\eta_K \geq \theta \max \eta_K$, for $0 < \theta < 1$, will be tagged and refined.

The BR error estimator for the interior condition is explicit and easily-computed. Ainsworth, Oden, and Lee [2] extend it to certain obstacle problems, but with heuristic aspects. On the other hand, for certain PDE problems, alternative estimators are provided by the *dual weighted residual* technique [6]. These are defined using a particular quantity of interest and some approximately-computed nonlocal weights, essentially Green's functions of the adjoint PDE. Suttmeier [32] has extended this weighted residual technique to certain VI problems, but it also requires heuristic steps even for the classical obstacle problem. In this paper we avoid the complexity of such techniques, and instead take a pragmatic approach which combines refinement near a computed free boundary with application of a PDE-type error estimator, as above, within the computed inactive sets.

4. New adaptive mesh refinement strategies. Our first two adaptive mesh refinement (AMR) methods, Algorithms 4.1 and 4.2, do targeted refinement near the free boundary. They are of *tag-and-refine* type. The third *metric-based mesh adaptation* approach [33], Algorithm 4.3, is more expensive.

These methods all start from a computed solution $u_h \in \mathcal{K}_h \subset \mathcal{X}_h$ to problem (2.7). Specifically we need only the mesh $\mathcal{T}_h$, the obstacle $\psi_h \in \mathcal{X}_h$, and u_h . The following two concepts are fundamental.

DEFINITION 4.1. Denote the vertices of $\mathcal{T}_h$ by x_j . Given $u_h \geq \psi_h$ and a tolerance $tol > 0$, the nodal active set indicator is the unique $\nu_h \in \mathcal{X}_h$ satisfying

$$\nu_h(x_j) = \begin{cases} 1, & u_h(x_j) - \psi_h(x_j) < tol \\ 0, & \text{otherwise.} \end{cases}$$

DEFINITION 4.2. An element marking of $\mathcal{T}_h$ is a piecewise-constant indicator function $\mathbb{1}_h \in DG_0(\mathcal{T}_h)$ with values in $\{0, 1\}$.

For PDE problems, and element marking can be derived from an error estimator η_K , as in the previous Section, associated with a quantity of interest such as energy or L^2 norm error [6]. However, a primary quantity of interest for obstacle problems is *vicinity to the unknown free boundary*. We do not associate this concept with a precise functional. Instead Algorithms 4.1 and 4.2 proceed heuristically to convert a nodal active set indicator into an element marking, with the goal of improving the approximation of the free boundary.

Queries of PETSc DMPlex objects [26] will also be needed, so let us sketch how this class supports unstructured meshes. A DMPlex object stores the topology (connectivity) of a dimension $d = 1, 2, 3$ mesh. Every mesh entity, regardless of dimension, is assigned a unique index. Mesh connectivity is understood as a stratified directed acyclic graph, where a covering/incidence relationship between mesh entities determines the graph edges. For example, a triangle within a 2D mesh is covered by its three edges, which are covered by their two endpoints (vertices). Each stratum ("height") in the DAG represents a dimensional class of mesh entity; for example the cells, edges, and vertices in a 2D mesh are at heights 0, 1, 2. Given the index p of a mesh entity, the basic DMPlex queries [5] are *cone*(p), the set of indices of entities which cover entity p , and its dual *support*(p), the set of entities which are covered by p . The transitive closure of *cone* is *closure*(p), and that of *support* is *star*(p). In our application we need only the vertices in the element closure; we denote this as $\text{closure}^\wedge(k)$. Similarly, $\text{star}_\vee(j)$ extracts only elements in the star of vertex x_j .

Now we can define the Unstructured Dilation Operator (UDO) method, Algorithm 4.1. It first computes a nodal active set indicator ν_h from u_h . Denoting the degrees of freedom in $\text{DG}_0(\mathcal{T}_h)$ by x_k , it then finds the set of indices k of elements such that $0 < \nu_h(x_k) < 1$. If $u_h \in \text{CG}_1$ then this condition holds when the element is incident to both active and inactive vertices. Then the method alternates closure and star on the marked element set n times, expanding the marking by n element layers, to generate a final element marking $\mathbb{1}_h$. The motivation here is that if the approximation $u_h \approx u$ is good for a non-degenerate VI problem (2.2) then the true free boundary Γ_u should pass through the elements indicated *a posteriori* by $\mathbb{1}_h$. In Sections 5 and 6 we will only consider $n = 1, 2$, as this much expansion seems to suffice for accurate representation of the free boundary after a few refinements.

Algorithm 4.1 Unstructured Dilation Operator (UDO) element marking

Require: mesh $\mathcal{T}_h$, solution $u_h \in \mathcal{K}_h$, obstacle $\psi_h \in \mathcal{X}_h$, tolerance $\text{tol} > 0$, and expansion parameter $n \geq 1$.

- 1: Compute nodal active set indicator $\nu_h \in \mathcal{X}_h$ for u_h .
- 2: Find initial element index set S_0 : $k \in S_0$ if $\nu_h(x_k) \in (0, 1)$.
- 3: **for** $i = 0, \dots, n - 1$ **do**

$$S_{i+1} = \bigcup_{k \in S_i} \bigcup_{j \in \text{closure}^\wedge(k)} \text{star}_\vee(j)$$

- 4: **return** marking $\mathbb{1}_h \in \text{DG}_0(\mathcal{T}_h)$ of all elements with indices in S_n .
-

While the UDO strategy explicitly manipulates indices, our second strategy, called Variable Coefficient Diffusion (VCD), Algorithm 4.2, is based on continuum ideas. Again the first step is to compute a nodal active set indicator $\nu_h \in \text{CG}_1$. This function is used as the initial condition of a time-dependent, variable-coefficient diffusion equation,

$$(4.1) \quad \frac{\partial s}{\partial t} = \nabla \cdot (D \nabla s), \quad s(t = 0) = \nu_h,$$

with Neumann (natural) boundary conditions, a well-posed problem. Clearly a solution of (4.1) at $t > 0$ is a smoothed form of ν_h . The diffusivity is set to the square of the element diameter h_K , namely $D = Ch_K^2 \in \text{DG}_0$, with $C = 0.5$ by default. The diffusion range is thus proportional to element diameter. In fact (4.1) is not solved exactly, or even very accurately, as we only compute $s_h \in \mathcal{X}_h$ from a single backward-Euler step of duration $\Delta t = 1$, equation (4.2) below. The default solver for this linear and elliptic problem is four iterations of conjugate gradient, preconditioned by incomplete-Cholesky factorization [11]. This inexpensive approximate solver has linear complexity in the vertices.

Figure 4 illustrates how the VCD algorithm applies to a one-dimensional obstacle problem. Note that the thresholds α and β are key parameters, with defaults $\alpha = 0.2$ and $\beta = 0.8$. Lowering α toward zero expands the marking further into the inactive set, away from the (computed) free boundary, while increasing β toward one expands further into the active set.

It turns out to be helpful, e.g. when solving the highly-nonlinear obstacle problem in Section 6, to add a minimum element diameter $h_{\min}$ to Algorithms 4.1 and 4.2. That is, after running one of the above Algorithms, we may un-set the element marking, $\mathbb{1}_h|_K = 0$, of an element with $h_K < h_{\min}$.

Algorithm 4.2 Variable Coefficient Diffusion (VCD) element marking

Require: mesh $\mathcal{T}_h$, solution $u_h \in \mathcal{K}_h$, obstacle $\psi_h \in \mathcal{X}_h$, tolerance $\text{tol} > 0$, and threshold interval $0 < \alpha < \beta < 1$.

- 1: Compute nodal active set indicator $\nu_h \in \mathcal{X}_h$ for u_h .
- 2: For h_K the element diameter, let $D = 0.5h_K^2 \in \text{DG}_0(\mathcal{T}_h)$.
- 3: Approximately solve for $s_h \in \mathcal{X}_h$, with natural boundary conditions:

$$(4.2) \quad s_h - \nabla \cdot (D \nabla s_h) = \nu_h$$

return marking $\mathbb{1}_h \in \text{DG}_0(\mathcal{T}_h)$ of all elements such that $s_h(x_k) \in (\alpha, \beta)$.




FIG. 4. Illustration of VCD: (a) Numerical solution u_h (solid), with nodes x_j shown (solid dots), and obstacle ψ_h (red dashed). (b) Nodal active set indicator $\nu_h \in \text{CG}_1(\mathcal{T}_h)$. (c) Smoothed indicator s_h , with element degrees of freedom x_k (circles) and thresholding levels (red). (d) Element marking $\mathbb{1}_h$; here 4 elements are marked for refinement.

From an element marking $\mathbb{1}_h$ we then apply *skeleton-based refinement* (SBR) [29] to generate a refined mesh. Elements with $\mathbb{1}_h = 1$ are refined, and any other elements as needed to avoid hanging nodes. Two SBR implementations are available, namely from PETSc DMPlex [5] and from the Netgen library [8]. Only the latter is currently capable of 3D refinement.

Our third AMR method, called *averaged-metric* (AVM; Algorithm 4.3), uses metric-based mesh adaptation [3]. In contrast to tag-and-refine methods, mesh adaptation generates a new mesh matching resolution and complexity targets. Adaptation is driven by an *a posteriori* metric field, defined as a continuous, matrix-valued function $M_h : \Omega \rightarrow \mathbb{R}^{d \times d}$ with each value $M_h(x)$ a symmetric and positive-definite matrix. Such a metric contains local information on distances, areas, and volumes [28]. From the metric the mesher itself generates a unit mesh [3]. In the original space the refined mesh has variable edge lengths and element aspect ratios.

AVM again starts from u_h and ψ_h . The first metric is isotropic, and it is computed from the gradient of a smoothed nodal active set indicator, namely s_h from the VCD method (Algorithm 4.2). The second metric is anisotropic, computed as an approximate Hessian of u_h via a Hessian-recovery technique [3]. (Actually the absolute value of the Hessian is used, which well-defined for a symmetric matrix [33].) These metrics are in the matrix-valued CG_1 FE space, with metric normalization constants

434 computed from target complexity and element diameter bounds in a standard man-
 435 ner [33]. By construction, the first metric should generate small elements near the
 436 free boundary, while the second should reduce FE approximation error in the inactive
 437 set, according to the standard interpolation theory [14]. In AVM the final metric
 438 is a weighted average of the two metrics. This final metric is both anisotropic and
 439 free-boundary focussed, and it implies refinement in both the active and inactive sets.
 440 Our implementation calls the Animate library (github.com/mesh-adaptation/animate)
 441 to construct, normalize, and average the metrics, and the Pragmatic library [19] for
 442 meshing at the last step. Note that steps 2 and 3 in Algorithm 4.3 use the target
 443 complexity and element diameter bounds parameters.

Algorithm 4.3 Averaged-metric (AVM) mesh adaptation

Require: mesh $\mathcal{T}_h$, solution $u_h \in \mathcal{K}_h$, obstacle $\psi_h \in \mathcal{X}_h$, target complexity N , ele-
 ment diameter bounds $0 < h_{\min} < h_{\max}$, and averaging weight $0 \leq \gamma \leq 1$

- 1: Compute $s_h \in \mathcal{X}_h$ from Algorithm 4.2, using u_h and ψ_h .
 - 2: Compute normalized isotropic free-boundary metric: $M_1(x) = c_1 |\nabla s_h(x)| I_{d \times d}$.
 - 3: Compute normalized anisotropic metric from Hessian: $M_2(x) = c_2 |H u_h(x)|$.
 - 4: Average the metrics: $M(x) = \gamma M_1(x) + (1 - \gamma) M_2(x)$.
 - 5: **return** new mesh $\tilde{\mathcal{T}}_h$ which is unit with respect to $M(x)$.
-

444 The right-hand image in Figure 3 (Introduction) shows an AVM result on the
 445 “ball” obstacle problem; see Figure 2, middle. Note that iterating the AVM method,
 446 even while holding the target mesh complexity constant, can be worthwhile because
 447 the increased resolution near the free boundary allows the *a posteriori* metric to be-
 448 come more effective. A key idea when iterating AVM is that cross-mesh interpolation
 449 [18] provides a high-quality initial iterate on the new mesh. Because of the nontrivial
 450 computations needed in metric-based methods [3, 33], at high resolution one AVM
 451 iteration is notably more expensive than an iteration of UDO or VCD.

452 All three Algorithms run in parallel under the MPI protocol used by Firedrake
 453 [26] and PETSc. However, only UDO produces results which are independent of the
 454 number of processes. The default preconditioned Krylov solver in VCD is slightly-
 455 dependent on process count [11]. Choosing a direct solver for problem (4.2) would
 456 give the VCD method process-independence, but it would also reduce scalability.

457 We end this Section with two Python examples which illustrate how to use the
 458 open source VIAMR library (github.com/StefanoFochesatto/viamr).

459 *Example 4.3.* Consider Example 1 from [2], a classical obstacle problem over a
 460 rectangle, with obstacle $\psi = 0$ and a known exact solution. The code below applies
 461 the Firedrake and VIAMR libraries to solve this problem. First it generates a uniform
 462 coarse mesh, then it applies VCD marking near the free boundary, and then it refines
 463 to a new mesh (Figure 5). There is no refinement in the inactive set, which is necessary
 464 for convergence; compare Example 4.4, and see the next Section.

```

1  from firedrake import *
2  from viamr import VIAMR
3
4  mesh = RectangleMesh(6, 12, 0.5, 1.0)
5  x, y = SpatialCoordinate(mesh)
6  r = (x + 1.0) ** 2 + y ** 2
7  uexact = conditional(r < 2.0, 0.25 * r - 0.5 - 0.5 * ln(0.5 * r), 0.0)
8
```

```

9  V = FunctionSpace(mesh, "CG", 1)
10 uh, vh = Function(V), TestFunction(V)
11 F = inner(grad(uh), grad(vh)) * dx - Constant(-1) * vh * dx
12 bcs = DirichletBC(V, Function(V).interpolate(uexact), "on_boundary")
13 problem = NonlinearVariationalProblem(F, uh, bcs)
14
15 sp = {"snes_type": "vinewtonrsls"}
16 solver = NonlinearVariationalSolver(problem, solver_parameters=sp)
17 psih = Function(V).interpolate(0.0)
18 INF = Function(V).interpolate(Constant(PETSc.INFINITY))
19 solver.solve(bounds=(psih, INF))
20
21 amr = VIAMR()
22 mark = amr.vcdmark(uh, psih)
23 VTKFile("mesh.pvd").write(uh, mark)
24
25 refinedmesh = amr.refinemarkedelements(mesh, mark)
26 VTKFile("refinedmesh.pvd").write(refinedmesh)

```




FIG. 5. The Python code in Example 4.3 solves Example 1 from [2]: initial mesh (left), marking by the default-parameter VCD scheme (middle), and the refined mesh (right). The exact free boundary, unknown to the algorithm, is overlaid in black.

465 *Example 4.4.* The following snippet was used for the left-hand mesh in Figure 3:

```

1  amr = VIAMR()
2  fbmark = amr.udomark(uh, psih, n=1)
3  residual = -div(grad(uh))
4  imark, _, _ = amr.brinactivemark(uh, psih, residual, theta=0.7)
5  mark = amr.unionmarks(fbmark, imark)
6  refinedmesh = amr.refinemarkedelements(mesh, mark)

```

466 Here the UDO method is used to mark near the free boundary, then the BR error
467 indicator (Section 3), from a computed element-wise residual, is applied for marking
468 in the inactive set. Finally the marks are unioned and a new mesh is generated.

469 **5. Results on classical obstacle problems.** We start this Section with a
470 revised look at how numerical errors should be measured for obstacle problems, and
471 then we show convergence and performance results on some examples. Section 6 will
472 show results for a more-challenging, non-classical obstacle problem.

In order to measure the quality of the approximate sets (2.8), versus the exact sets (1.2), and to determine rates of geometric convergence, the VIAMR library computes a distance between sets. Our implementation of (5.1) below is based on Firedrake's supermesh concept [18], so it is computable for sets S, T which are unions of elements from different meshes, and/or for sets defined by algebraic (conditional) expressions in the mesh coordinates.

DEFINITION 5.1. *The Jaccard distance [27] between measurable sets $S, T \subset \Omega$ is*

$$(5.1) \quad d(S, T) = 1 - \frac{|S \cap T|}{|S \cup T|},$$

where $|\cdot|$ is Lebesgue measure, with $d(S, T) = 0$ by definition if $|S \cup T| = 0$.

We will use Jaccard distance to compare active sets. If $d(S, T) = 0$ then S and T geometrically agree up to a set of measure zero, and for a computed active set A_u^h , defined by an element marking (Definition 4.2), we will report $d(A_u^h, A_u)$ when A_u is exactly known.

A second goal for numerical solutions of obstacle problems is to avoid wasted effort in the active set, once the free boundary is well-resolved. As already discussed, if $\Gamma_u^h \approx \Gamma_u$ is a good approximation then the obstacle data itself, namely ψ , can be used to represent the solution in the computed active set A_u^h . This motivates a definition.

DEFINITION 5.2. *Suppose the original VI problem (2.2) has exact obstacle $\psi \in \mathcal{X} \cap C(\bar{\Omega})$. Consider a computed solution u_h to the approximating problem (2.7), with computed active set A_u^h , a union of closed elements. We define the preferred approximation $\tilde{u}_h$ as the measurable function*

$$(5.2) \quad \tilde{u}_h(x) = \begin{cases} \psi(x), & x \in A_u^h \\ u_h(x), & \text{otherwise.} \end{cases}$$

For a classical obstacle problem (1.4), the preferred approximation $\tilde{u}_h$ is in $L^2(\Omega)$. It is generally discontinuous, even when $\mathcal{X}_h = \text{CG}_k$, thus $\tilde{u}_h \notin H^1(\Omega)$. However, $\tilde{u}_h$ is defined *without* reference to the exact solution to problem (2.2); only the continuum data ψ is referenced. The L^2 error relative to $\tilde{u}_h$ has a clear decomposition over the sets defined in (1.2) and (2.8):

$$(5.3) \quad \|u - \tilde{u}_h\|_2^2 = 0 + \int_{A_u \setminus A_u^h} |\psi - u_h|^2 + \int_{A_u^h \setminus A_u} |u - \psi|^2 + \int_{I_u \cap I_u^h} |u - u_h|^2.$$

The leading zero comes from the integral over $A_u \cap A_u^h$, where $u = \tilde{u}_h = \psi$. The next two terms are small if the free boundary has been accurately located ($\Gamma_u^h \approx \Gamma_u$), the goal of AMR in Section 4. The final term is small if inactive-set refinement is effective (Section 3). The next example shows that our methods converge in norm (5.3).

Example 5.3. Consider the “ball” classical obstacle problem shown in Figure 1 and Figure 2 (middle), which has a known exact solution [11, Chapter 12]. Algorithms 4.1 (UDO) and 4.2 (VCD) were applied to this problem, combined with inactive-set marking based on the BR error estimator (3.3). These methods, denoted UDO+BR and VCD+BR, avoid refinement in the active set, except near the estimated free boundary. We also applied mesh adaptation Algorithm 4.3 (AVM), setting mesh complexity targets to closely-match the number of elements from the other Algorithms; note that

AVM refines everywhere. The experiment started from a coarse uniform mesh over the square domain, and applied 11 levels of AMR, so that the number of elements increased by four orders of magnitude. We also compared uniform refinement.

As shown in Figure 6, our AMR methods are clearly superior to uniform refinement when evaluated by active set Jaccard distances $d(A_u^h, A_u)$. At higher resolutions, all three AMR methods reduce this distance an order of magnitude below uniform refinement. While AVM produces the best meshes by this measure, its runtime is about an order of magnitude higher at comparable mesh complexity (not shown).




FIG. 6. Active set Jaccard distances $d(A_u^h, A_u)$.

On the other hand, Figure 7 (left) shows that L^2 convergence using the standard error norm $\|u - u_h\|_2$ stagnates for the UDO+BR method. This is entirely due to the contribution from the active set, i.e. from $|\psi(x) - \psi_h(x)|$ for $x \in A_u^h$, over meshes which have *deliberately* not been refined in A_u^h . If we compute the L^2 error relative to the preferred approximation for the same meshes, namely $\|u - \tilde{u}_h\|_2$ as in (5.3), then the convergence is better than from uniform refinement, in a per-element sense. Results from VCD+BR are virtually identical, and thus not shown.

As shown in Figure 7 (right) for the AVM algorithm, because of its active set refinement, standard ($\|u - u_h\|_2$) and preferred ($\|u - \tilde{u}_h\|_2$) error norms are comparable.




FIG. 7. Left: L^2 norm errors for the UDO+BR method, versus uniform refinement. Right: The same norm errors for the AVM method.

To illustrate some additional mesh and performance features of our AMR ap-

proaches, we consider three more classical obstacle problem examples. These do not have known exact solutions.

Example 5.4. Consider the spiral example on $\Omega = (-1, 1)^2$ from [20, subsection 7.1.1], shown in Figure 2 (left). This example has source term $f = 0$, homogeneous boundary values $g = 0$, and a nontrivial obstacle ψ . The active set is small in area while the free boundary is long. A result from applying the UDO+BR approach is shown in Figure 8. (Results from VCD+BR and AVM approaches are very similar and not shown.) The free boundary is well-resolved, but there is no performance advantage from avoiding refinement in such a small active set.




FIG. 8. A refined mesh for Example 5.4, with 2×10^5 elements, from seven levels of the UDO+BR approach, starting from an initial uniform mesh of 200 elements. Element diameter (resolution) is $h \approx 10^{-3}$ along the free boundary.

Example 5.5. Consider the example shown in Figure 2 (right), which has a large active set. Here $\Omega = (0, 1)^2$ and $\psi = g = 0$ in (1.4). The source term f is a sum of small-variance gaussian peaks against a negative background value; see the code for details. The blistering property (Appendix A) applies, so inactive set components necessarily include a $f > 0$ portion. The inactive set has six connected components, but this is only revealed at high resolution. Figure 9 shows the result of applying the VCD+BR method, giving a mesh with 100 times fewer elements than a uniform mesh with the same free-boundary resolution. Similar to the glacier example in the next Section, in this case the avoidance of active-set refinement gives a distinct performance advantage. The result from UDO+BR is very similar, and not shown.

Example 5.6. Our third example is a free-boundary variation of the classic L-shaped Laplace equation problem with an interior corner. Here the obstacle ψ is the upper unit hemisphere, centered at the origin, and the domain is $\Omega = (-2, 5)^2 \setminus S$ where S is a 3×3 square in the lower right corner. Since $f = 0$ and $g = -1$, the Laplace equation is solved in the inactive set, with unbounded solution Hessian in the interior corner. A result from AVM, starting from a coarse unstructured mesh, is shown in Figure 10. (Again the other methods produce similar results.) Though the exact free boundary is not known, it resolves into a smooth and nearly-circular curve. At the same time there is refinement in the interior corner. This example shows




FIG. 9. Left: A mesh for Example 5.5 of 4.7×10^5 elements, with resolution $h \approx 10^{-3}$ along the free boundary, from four levels of the VCD+BR approach, starting from a uniform mesh of 1800 elements. Right: Zooming-in reveals inactive-set refinement, and separation of components.

559 that free-boundary localization and classical PDE refinement goals are simultaneously
560 addressed.




FIG. 10. A refined mesh from three iterations of AVM Algorithm 4.3 on Example 5.6, showing a well-resolved free boundary and refinement in the interior corner.

561 **6. Application to determining glaciated land areas.** In this section we
562 apply our AMR techniques to a model for the steady geometry of a glacier, over a
563 given bedrock topography and subject to a given climate. In particular, this model
564 solves for which land is covered by ice. Though glacier ice will cover any land on which
565 snowfall exceeds melt, ice flow expands the glaciated area out to a free boundary, the
566 glacier margin, which can only be found from conservation equations. In fact, any
567 fluid layer subject to surface processes which add or remove mass is governed by a
568 related model [10]. In these VI problems, which generally are *not* of optimization

type, the unknown fluid thickness must be nonnegative, the operator is nonlinear, and the blistering property (Appendix A) applies.

Let $\Omega \subset \mathbb{R}^2$ be a fixed land region, with bedrock elevation $b \in C^1(\Omega)$ and a surface mass-balance [21] function $a : \Omega \times \mathbb{R} \rightarrow \mathbb{R}$. The value of this climatic function is the annually-averaged rate of ice accumulation (snow) minus melt and runoff. If a depends on the surface elevation s then we suppose that $a(x, y, s) \in L^\infty(\Omega)$ for any $s \in L^\infty(\Omega)$, and additionally that a is Lipschitz continuous in s .

We use an ice flow approximation called the *shallow ice approximation* [21], derived by a small aspect ratio argument from mass and momentum conservation, in its simplest isothermal and non-sliding case. Applying the standard exponent for the shear-thinning flow law of ice [21], we have a VI for admissible transformed ice thickness functions $u \in \mathcal{K} = \{u \in \mathcal{X} : u \geq 0 \text{ and } u|_{\partial\Omega} = 0\}$, in the Banach space $\mathcal{X} = W^{1,4}(\Omega)$ [22]. The thickness itself is $H = u^{3/8}$, while the surface elevation is $s = H + b = u^{3/8} + b$. The model is a VI based upon a “tilted” variation of the p -Laplacian operator [22]:

$$(6.1) \quad \int_{\Omega} \Gamma |\nabla u + \beta(u)|^2 (\nabla u + \beta(u)) \cdot \nabla(v - u) - \tilde{a}(u)(v - u) \, dx \geq 0$$

for all $v \in \mathcal{K}$. Here $\Gamma > 0$ is an ice softness constant, the tilt $\beta(u) = \frac{8}{3}u^{5/8}\nabla b$ is a nonlinear multiple of the bed gradient, and $\tilde{a}(u) = a(x, y, u^{3/8} + b)$ denotes the surface mass balance. If a has no s dependence then we may treat it as the source term, and write “ $F(u)[v - u] \geq \ell[v - u]$ ” as in (2.2), but otherwise we simply set $\ell = 0$ in (2.2). If also $\nabla b = 0$ then (6.1) defines a nonlinear operator $F(v)[w]$ which is 4-coercive; see Example 2.2. Regarding the mathematical theory of (6.1), existence holds for any b when a is independent of s [22]. Simple cases of *non*-existence are known when a depends on s [23], but our examples below have slow increase of snowfall with elevation [21], which avoids this issue, and we suppose the solutions below are unique.

The strong-form interior condition for (6.1) is a conservation equation, namely $\nabla \cdot \mathbf{q} = \tilde{a}$ over I_u , where the flux is given by

$$(6.2) \quad \mathbf{q} = -\Gamma |\nabla u + \beta(u)|^2 (\nabla u + \beta(u)).$$

At the free boundary (glacier margin) the solution u , the solution gradient $|\nabla u|$, and the flux $\mathbf{q}$ all go to zero. On the other hand, observations confirm that the gradient of a glacier’s surface has large magnitude as one approaches the margin from the ice side, a property which emerges from (6.1) when one differentiates the associated surface elevation $s = u^{3/8} + b$. Generally $|\nabla s| \rightarrow \infty$ at the glacier margin.

We apply a conforming, continuous, piecewise-linear FE method to (6.1), over unstructured triangulations. The FE problem is either solved directly by the same VI-adapted Newton solver used in Section 5, namely `vinewtonrs1s` in PETSc [5], or with an added outer Picard-type iteration over the tilt $\beta(u)$ [22] and source $\tilde{a}(u)$, with Newton as the inner iteration. The Newton steps are solved using Firedrake’s default direct linear solver. Very similar results are obtained when these iterations both converge, but the Picard-based solver is more robust when $|\nabla b|$ is large or $\tilde{a}$ depends on u .

We consider two particular problems, each over a square domain $\Omega = (0, L)^2$ with $L = 1800$ kilometers. The first “dome” problem has a flat bed, elevation-independent source, and a known exact solution [9]; this is used for verification and to compare AMR algorithms. The second problem has a bumpy bed and an elevation-dependent surface mass balance function.

Example 6.1. The dome problem was solved with 13 levels of refinement from an initial, uniform $h_0 = 360$ kilometer mesh to a final mesh with glacier margin resolution of $h_{13} \approx 30$ meters. Both UDO and VCD methods (Algorithms 4.1 and 4.2) were applied, with gradient recovery (GR; Example 3.1) marking in the inactive set,¹ and the results were very similar. Because the obstacle is $\psi = 0$, the preferred (5.2) and standard numerical solutions are the same. Figure 11 shows relative H^1 errors in the solution ($\|u - u_h\|_{H^1} / \|u\|_{H^1}$) and absolute L^∞ errors in the corresponding ice thickness ($\|H - H_h\|_{L^\infty}$), versus number of elements. The latter errors are much more responsive to AMR because the ice thickness gradient is singular at the free boundary. That is, Figure 11 right is primarily reporting (lateral) margin location errors. In fact, similar to the Jaccard distances reported in Example 5.3, Figure 12 shows the maximum radial error for the numerical margin. We see that our AMR methods are quite efficient for free-boundary localization, compared to uniform refinement.




FIG. 11. Left: Relative H^1 norm errors in u are comparable for AMR and uniform refinement, in a per-element sense. Right: AMR is very effective for the maximum thickness error.




FIG. 12. AMR generates accurate free-boundary locations. To achieve tens of meters accuracy, as here, uniform refinement would use orders of magnitude more elements.

Example 6.2. The next example is more realistic. It uses a bumpy bedrock topography $b(x, y)$, constructed from a finite sum of sinusoids [13, Example 8.4], and a surface mass balance function $a(s)$ which depends only on surface elevation; see the

¹Note that the argument for BR marking given in Section 3 does not apply here.

631 source code for details. An important parameter in the formula for $a(s)$ is the equi-
 632 librium line altitude (ELA) [21], denoted s_0 , with $a > 0$ above this altitude and $a < 0$
 633 below. To illustrate this parameter, Figure 13 shows the topography along a transect
 634 $x = \hat{x}$, the red line in Figure 14, with three s_0 values superimposed. Net snowfall
 635 $a > 0$ occurs on a solution-dependent set $S = \{(x, y) : u(x, y)^{3/8} + b(x, y) > s_0\}$, and
 636 glaciation then extends beyond S because of flow: $I_u \supset S$. The difficult nonlinearities
 637 in model (6.1) include this elevation-dependent surface mass balance, but also the
 638 $u^{5/8}$ dependence in the tilt $\beta(u)$. Among the effects of the latter is the tendency of
 639 ice to pool in bedrock elevation lows. Our primary goal in this example is to generate
 640 an accurate map of glaciation, i.e. of the inactive set I_u . Related goals are to compute
 641 the ice volume $V = \int_{\Omega} u^{3/8} dx dy$, and to find the number of connected components
 642 of I_u , that is, the number of glaciers in the given climate.




FIG. 13. Bedrock elevation (solid) along a transect, with three ELA levels s_0 .

643 We again applied UDO, with $n = 2$ expansion, plus GR marking in the inactive
 644 set to resolve ice flow. Using a minimum element diameter $h_{\min} = 500$ meters, the
 645 final meshes averaged 300 meter resolution along most of their margins. The solution
 646 surface elevation maps for the three ELA values are shown in Figure 14. The very
 647 different ice volumes, 1.1×10^5 , 5.6×10^5 , and 8.9×10^5 cubic kilometers, respectively,
 648 reflect the strong ELA sensitivity of realistic glacier models [21].²




FIG. 14. Glacier surface elevation (grayscale) for ELA of 1000 meters (left), 800 meters (middle), and 600 meters (right). Inset boxes are 200 km across; see Figure 15.

649 Figure 15 shows details of the meshes along sample portions of the glacier margins.
 650 It is key to observe that under the UDO+GR approach the mesh was *not* refined in the
 651 vast majority of the active set, and this represents a significant efficiency. That is, we
 652 did not waste computational effort (e.g. elements) on modeling ice flow in unglaciated
 653 areas. In fact this highly-nonlinear problem causes this AMR algorithm to work quite
 654 hard to pin down the glacier margin, as only once the flow on the inactive side of the
 655 margin is resolved at high resolution will margin location stabilize.

²Compare the volume of the present-day Greenland ice sheet at 2.9×10^6 cubic kilometers.




FIG. 15. Details of refined meshes (red) along glacier margins, from the inset boxes in Figure 14, for the same ELA values. The dark areas are coarsely-meshed ice-free areas, i.e. active sets.

7. Discussion. It is well-known that inequality-constrained obstacle problems are intrinsically nonlinear, thus that solution methods must be iterative. At a basic level our AMR methods simply extend the mandatory iteration to the generation of sequences of well-adapted unstructured meshes.

The UDO and VCD Algorithms from Section 4 combine two relatively-simple tagging strategies based *a posteriori* on the numerical solution:

- i) assuming u_h and ψ_h can estimate the true free boundary Γ_u , mark nearby elements on both sides of the computed free boundary Γ_u^h , and
- ii) assuming u_h approximates the true solution u in the computed inactive set I_u^h , mark elements in this set using PDE-type error estimates (Section 3).

AVM Algorithm 4.3 uses the same essential heuristics to generate a metric, which enables mesh adaptation.

Our AMR strategies are tied to the *a priori* theory in Section 2. On the other hand, we have no direct, quantitative theory of how they reduce numerical error. Observe that i) has nothing directly to do with the regularity or residual of u_h ; refinement is geometrical in the domain. On the other hand, ii) could be done by methods of arbitrary sophistication. For example, dual weighted residual methods [6, 32] could be applied, although their adjoint-type weights are incapable of communicating information between disconnected components of the inactive set, the number and geometry of which are only found at solution time. (Example 5.5 is an example.)

For classical obstacle problems, the mesh sequences built by our AMR techniques allow the bulk of the Newton iterations to occur inexpensively on the coarse early meshes. However, for harder problems like the glaciation obstacle problem in Section 6, the solver and AMR components strongly interact. Nontrivial solver iterations are needed on intermediate meshes, to find and stabilize the free boundary.

An important extension of our work would be to find effective refinement strategies for time-dependent obstacle problems and parabolic VIs. Also, performance could be improved by integrating AMR into a multilevel solver, though this is also for future research. For example, the solver in [13] uses coarse meshes to make large corrections, including geometrically to the free boundary, and so a combination with AMR promises highly-scalable solutions of difficult obstacle problems.

REFERENCES

- [1] M. AINSWORTH AND J. T. ODEN, *A Posteriori Error Estimation in Finite Element Analysis*, John Wiley & Sons, Inc., New York, 2000.

- [2] M. AINSWORTH, J. T. ODEN, AND C. LEE, *Local a posteriori error estimators for variational inequalities*, Numerical Methods for Partial Differential Equations, 9 (1993), pp. 23–33.
- [3] F. ALAUZET, *Metric-Based Anisotropic Mesh Adaptation*, 2010, https://pages.saclay.inria.fr/frederic.alauzet/cours/cea2010_V3.pdf. Course material, CEA-EDF-INRIA Numerical Analysis Summer School.
- [4] I. BABUŠKA AND W. RHEINBOLDT, *On the reliability and optimality of the finite element method*, Computers & Structures, 10 (1979), pp. 87–94, [https://doi.org/10.1016/0045-7949\(79\)90076-2](https://doi.org/10.1016/0045-7949(79)90076-2).
- [5] S. BALAY, S. ABHYANKAR, M. F. ADAMS, S. BENSON, J. BROWN, P. BRUNE, K. BUSCHELMAN, E. CONSTANTINESCU, L. DALCIN, A. DENER, V. EIJKHOUT, J. FAIBUSSOWITSCH, W. D. GROPP, V. HAPLA, T. ISAAC, P. JOLIVET, D. KARPEEV, D. KAUSHIK, M. G. KNEPLEY, F. KONG, S. KRUGER, D. A. MAY, L. C. MCINNES, R. T. MILLS, L. MITCHELL, T. MUNSON, J. E. ROMAN, K. RUPP, P. SANAN, J. SARICH, B. F. SMITH, H. SUH, S. ZAMPINI, H. ZHANG, H. ZHANG, AND J. ZHANG, *PETSc/TAO User's Manual*, Tech. Report ANL-21/39 - Revision 3.22, Argonne National Laboratory, 2025, <https://doi.org/10.2172/2205494>.
- [6] W. BANGERTH AND R. RANNACHER, *Adaptive Finite Element Methods for Differential Equations*, Lectures in Mathematics ETH Zürich, Birkhäuser, Basel Berlin, 2003.
- [7] S. BENSON AND T. MUNSON, *Flexible complementarity solvers for large-scale applications*, Optimization Methods and Software, 21 (2006), pp. 155–168.
- [8] J. BETTERIDGE, P. E. FARRELL, M. HOCHSTEGER, C. LACKNER, J. SCHÖBERL, S. ZAMPINI, AND U. ZERBINATI, *ngsPETSc: A coupling between NETGEN/NGSolve and PETSc*, Journal of Open Source Software, 9 (2024).
- [9] E. BUELER, *Stable finite volume element schemes for the shallow ice approximation*, J. Glaciol., 62 (2016), pp. 230–242.
- [10] E. BUELER, *Conservation laws for free-boundary fluid layers*, SIAM J. Appl. Math., 81 (2021), pp. 2007–2032, <https://doi.org/10.1137/20M135217X>.
- [11] E. BUELER, *PETSc for Partial Differential Equations: Numerical Solutions in C and Python*, no. 31 in Software, Environments, and Tools, SIAM Press, Philadelphia, 2021.
- [12] E. BUELER, *Surface elevation errors in finite element Stokes models for glacier evolution*, 2024, <https://arxiv.org/abs/2408.06470v2>.
- [13] E. BUELER AND P. E. FARRELL, *A full approximation scheme multilevel method for nonlinear variational inequalities*, SIAM J. Sci. Comput., 46 (2024), pp. A2421–A2444, <https://doi.org/10.1137/23m1594200>.
- [14] P. CIARLET, *The Finite Element Method for Elliptic Problems*, SIAM Press, Philadelphia, 2002. Reprint of the 1978 original.
- [15] H. C. ELMAN, D. J. SILVESTER, AND A. J. WATHEN, *Finite Elements and Fast Iterative Solvers: with Applications in Incompressible Fluid Dynamics*, Numerical Mathematics and Scientific Computation, Oxford University Press, 2nd ed., 2014.
- [16] L. C. EVANS, *Partial Differential Equations*, Graduate Studies in Mathematics, American Mathematical Society, Providence, 2nd ed., 2010.
- [17] R. S. FALK, *Error estimates for the approximation of a class of variational inequalities*, Mathematics of Computation, 28 (1974), pp. 963–971.
- [18] P. FARRELL, M. PIGGOTT, C. PAIN, G. GORMAN, AND C. WILSON, *Conservative interpolation between unstructured meshes via supermesh construction*, Computer Methods in Applied Mechanics and Engineering, 198 (2009), pp. 2632–2642, <https://doi.org/10.1016/j.cma.2009.03.004>.
- [19] G. GORMAN, J. SOUTHERN, P. FARRELL, M. PIGGOTT, G. ROKOS, AND P. KELLY, *Hybrid openmp/mpi anisotropic mesh smoothing*, in Proceedings of the International Conference on Computational Science, vol. 9, 2012, pp. 1513–1522, <https://doi.org/10.1016/j.procs.2012.04.166>.
- [20] C. GRÄSER AND R. KORNUBER, *Multigrid methods for obstacle problems*, Journal of Computational Mathematics, 27 (2009), pp. 1–44, http://global-sci.org/intro/article_detail/jcm/8559.html.
- [21] R. GREVE AND H. BLATTER, *Dynamics of Ice Sheets and Glaciers*, Advances in Geophysical and Environmental Mechanics and Mathematics, Springer, Berlin, Germany, 2009.
- [22] G. JOUVET AND E. BUELER, *Steady, shallow ice sheets as obstacle problems: well-posedness and finite element approximation*, SIAM J. Appl. Math., 72 (2012), pp. 1292–1314.
- [23] G. JOUVET, J. RAPPAZ, E. BUELER, AND H. BLATTER, *Existence and stability of steady-state solutions of the shallow-ice-sheet equation by an energy-minimization approach*, J. Glaciol., 57 (2011), pp. 345–354.
- [24] B. KEITH AND T. M. SUROWIEC, *Proximal Galerkin: A structure-preserving finite ele-*

- ment method for pointwise bound constraints, *Foundations of Computational Mathematics*, (2024), pp. 1–97, <https://doi.org/10.1007/s10208-024-09681-8>.
- [25] D. KINDERLEHRER AND G. STAMPACCHIA, *An Introduction to Variational Inequalities and their Applications*, Academic Press, New York, 1980.
- [26] M. LANGE, L. MITCHELL, M. G. KNEPLEY, AND G. J. GORMAN, *Efficient mesh management in Firedrake using PETSc DMPlex*, *SIAM Journal on Scientific Computing*, 38 (2016), pp. S143–S155.
- [27] M. LEVANDOWSKY AND D. WINTER, *Distance between sets*, *Nature*, 234 (1971), pp. 34–35, <https://doi.org/10.1038/234034a0>.
- [28] A. LOSEILLE AND F. ALAUZET, *Continuous mesh framework part I: well-posed continuous interpolation error*, *SIAM Journal on Numerical Analysis*, 49 (2011), pp. 38–60.
- [29] A. PLAZA AND G. F. CAREY, *Local refinement of simplicial grids based on the skeleton*, *Applied Numerical Mathematics*, 32 (2000), pp. 195–218.
- [30] W. K. PRATT, *Digital Image Processing*, John Wiley & Sons, Inc., 1991.
- [31] 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 (2016), pp. 24:1–24:27.
- [32] F.-T. SUTTMER, *Numerical Solution of Variational Inequalities by Adaptive Finite Elements*, Vieweg + Teubner, 2008.
- [33] J. G. WALLWORK, N. BARRAL, D. A. HAM, AND M. D. PIGGOTT, *Anisotropic goal-oriented mesh adaptation in Firedrake*, in *Proceedings of the 28th International Meshing Roundtable*, S. Shontz, J. Peiró, and R. Viertel, eds., Buffalo, NY, 2020, pp. 83–100, <https://doi.org/10.5281/zenodo>.

Appendix A. Obstacle problems with the blistering property.

In a classical, unilateral obstacle problem (1.4) an upward force $f(x) > 0$ may occur in the active set A_u . Physically speaking, in such a case the upward force was insufficient to lift u off the obstacle ψ . However, in certain problems this situation cannot occur. The hypotheses of Lemma A.1 hold for classical problems (1.4) when $\psi = 0$, for the porous dam saturation free-boundary problem [2, for example], and for the glacier problem in Section 6.

LEMMA A.1. *Suppose that $\psi = 0$ and $\ell[v] = \int_{\Omega} f v \, dx$ for $f \in L^2(\Omega)$. Suppose also that the operator F is given by measurable densities,*

$$(A.1) \quad F(w)[v] = \int_{\Omega} \phi(w(x), \nabla w(x)) v(x) + \Phi(w(x), \nabla w(x)) \cdot \nabla v(x) \, dx,$$

satisfying $\phi(0, 0) = 0$ and $\Phi(0, 0) = 0$. Then for u solving VI (2.2), $f \leq 0$ a.e. in A_u .

Proof. Let $S \subset A_u$ be a Borel set, and denote its indicator function by $\chi_S \in L^\infty(\Omega)$. By a density argument and the assumptions on ϕ and Φ , the operator value $F(u)[\chi_S]$ is well-defined, and zero, because S is in the active set. Thus by Lemma 2.3,

$$0 \leq d\mu_u(S) = F(u)[\chi_S] - \ell[\chi_S] = 0 - \ell[\chi_S] = - \int_S f \, dx.$$

This shows $f \leq 0$ a.e. with respect to Lebesgue measure. $\square$

DEFINITION A.2. *For a unilateral obstacle problem (2.2) with source term $\ell[v] = \int_{\Omega} f v \, dx$ and $f \in C(\bar{\Omega})$, we say that the blistering property holds if $A_u \subset \{x \in \Omega : f(x) \leq 0\}$ [22], that is, if the conclusion of Lemma A.1 holds.*

To explain the language, if $f(x) > 0$ for some $x \in \Omega$ then, assuming this property, the inactive set I_u must be non-empty, and $x \in I_u$. That is, even a small upward force “blisters” the membrane off the obstacle.

Any classical problem (1.4) with a smooth obstacle $\psi \in C^2(\bar{\Omega})$ can be transformed into one with the blistering property. Let $\tilde{v} = v - \psi$, thus $\tilde{\mathcal{K}} = \{\tilde{v} \in \mathcal{X} : \tilde{v} \geq$

800 0 and $\tilde{v}|_{\partial\Omega} = g - \psi|_{\partial\Omega}$. Integration-by-parts shows (1.4) is equivalent to finding
 801 $\tilde{u} = u - \psi \in \tilde{\mathcal{K}}$ so that

$$802 \quad (\text{A.2}) \quad \int_{\Omega} \nabla \tilde{u} \cdot \nabla (\tilde{v} - \tilde{u}) \geq \int_{\Omega} (f + \nabla^2 \psi) (\tilde{v} - \tilde{u}) \quad \text{for all } \tilde{v} \in \tilde{\mathcal{K}}.$$

803 For VI problem (A.2), if $x \in A_{\tilde{u}} = \{\tilde{u}(x) = 0\}$ then $\tilde{f}(x) = f(x) + \nabla^2 \psi(x) \leq 0$.

804 For blistering-property problems, AMR costs can be reduced by systematically
 805 avoiding refinement in the computed active sets. From a computed solution u_h , one
 806 identifies elements $K \in \mathcal{T}_h$ such that $K \subset A_u^h$. (In our implementation we require that
 807 every node in the closure of K is active, according to some tolerance.) The degenerate
 808 case must be excluded, so we also require $K \subset \Omega_- = \{x \in \Omega : f(x) < 0\}$; note this
 809 uses only the data of the problem. Then for $K \subset \Omega_- \cap A_u^h$ no refinement is needed
 810 if K is geometrically far from the free boundary. However, if later refinements move
 811 the free boundary to be incident to K then refinement of K becomes appropriate.