

Sharp inf-sup estimate for the Stokes equation in tight domains with periodic pillars and some numerical implications

Qi Xin^{a,b}, Shihua Gong^{a,b,c,*}, Jinchao Xu^d

^a*School of Science and Engineering, The Chinese University of Hong Kong, Shenzhen, 518172, Guangdong, China*

^b*Shenzhen International Center for Industrial and Applied Mathematics, Shenzhen Research Institute of Big Data, Shenzhen, 518000, Guangdong, China*

^c*Shenzhen Loop Area Institute, Shenzhen, 518000, Guangdong, China*

^d*Applied Mathematics and Computational Sciences, CEMSE Division, King Abdullah University of Science and Technology, Thuwal, 23955, Saudi Arabia*

Abstract

The predictive simulation of fluid dynamics in densely packed microfluidic devices, such as Deterministic Lateral Displacement (DLD) arrays, stagnates with standard iterative solvers. We show that this failure is not algorithmic but rooted in the pre-asymptotic degradation of the pressure-velocity coupling stability. For periodic pillar geometries in a generalized lattice framework, we prove that the continuous Ladyzhenskaya-Babuška-Brezzi (LBB) condition, also called the inf-sup constant, deteriorates exactly as m^{-1} up to a positive multiplicative constant, where m is the pillar density (the number of pillars per unit length). This induces a priori error amplification proportional to m and a pressure Schur complement condition number scaling as $\mathcal{O}(m^2)$. To overcome this theoretical limit, we propose a parameter-free, adaptively scaled Augmented Lagrangian (AL) stabilization strategy with penalty $\gamma \propto m^2$. Numerical experiments on both standard square and asymmetric DLD arrays validate the theoretical bounds: the AL method reduces outer FGMRES iterations from 437 to 22 on a 1.85M-DoF square array and from 687 to 24 on a 1.77M-DoF DLD array.

Keywords: Stokes equation, inf-sup estimate, finite element methods, preconditioning

1. Introduction

Microfluidic devices featuring dense periodic pillar arrays have advanced biomedical engineering, chemical processing, and precise particle manipulation. A prominent exemplar is the Deterministic Lateral Displacement (DLD) device [1, 2, 3, 4], widely employed in high-throughput cell sorting and biophysical fractionations. These devices typically consist of thousands to millions of microscopic solid pillars arranged in a shifted Bravais lattice to control fluid streamlines and particle trajectories.

While the fabrication of such devices has matured, the predictive computational modeling of the fluid dynamics and fluid-structure interactions (FSI) within these dense arrays remains open. As the pillar density increases (quantified by the number $m \rightarrow \infty$ of pillars in a unit length, or the characteristic cell size $\epsilon \rightarrow 0$), standard numerical solvers for the incompressible Stokes equations stagnate [5].

The performance of standard iterative frameworks and classical block-preconditioning techniques (e.g., standard Schur complement approximations) degrades in dense pillar arrays. The root cause lies in the mathematical structure of the saddle-point problem. The convergence rate of preconditioned iterative solvers is dictated by the stability of the pressure-velocity coupling, quantified by the continuous and discrete inf-sup (LBB) constants [6, 7, 8]. The inf-sup constant depends on the underlying domain geometry, deteriorating in elongated, anisotropic, or heavily perforated configurations [9, 10, 11]. As the pillar packing density increases, the stability of the system collapses, rendering standard parameter tuning ineffective.

The magnitude of the inf-sup constant β , see definition in (7), dictates both the continuous physical stability and the numerical resolvability of the fluid system:

- **A Priori Error Amplification:** At the level of numerical discretization, Brezzi's theorem [12, 6, 13, 14] establishes that the a priori error estimate for any stable mixed finite element method is amplified

*Corresponding author

Email addresses: qixin1@link.cuhk.edu.cn (Qi Xin), gongshihua@cuhk.edu.cn (Shihua Gong), jinchao.xu@kaust.edu.sa (Jinchao Xu)

by a factor proportional to β_h^{-1} , where β_h is the discrete inf-sup constant. If the continuous constant β approaches zero, the discrete inf-sup constant β_h inherits this geometric degradation [6], guaranteeing a loss of accuracy. As the geometric penalty drives β_h toward zero, this pre-factor amplifies the discretization error. The optimal asymptotic rates are overwhelmed unless prohibitively fine meshes offset this algebraic amplification.

- **Stagnation of Iterative Solvers:** At the linear algebra level, the condition number of the pressure Schur complement operator S is proportional to β_h^{-2} . This ill-conditioning dictates the stagnation of classical segregated iterative methods (such as Uzawa or block preconditioner algorithms) [7, 15] and neutralizes the efficiency of standard block-diagonal preconditioners in Krylov subspace methods (e.g., GMRES).

While the degradation of the continuous inf-sup constant in slender domains is documented in the literature [9, 10, 11], its precise pre-asymptotic estimate within the multiply-connected topologies of dense pillar arrays lacks explicit quantitative characterization. Classical homogenization theory [16, 17] studies the asymptotic limit as the characteristic cell size $\epsilon \rightarrow 0$ (or equivalently, the pillar density $m \rightarrow \infty$), smearing out the discrete solids to yield a continuous macroscopic Darcy’s law. However, practical CFD simulations of DLD devices operate in the pre-asymptotic regime: the pillars are discrete and finite, and the full Stokes equations must be solved explicitly to resolve local shear stresses for precise particle tracking.

We establish the sharp pre-asymptotic degradation rate of the continuous LBB constant β in these generalized periodic pillar geometries. By constructing global linear pressure fields and leveraging local restriction operators, we prove that the inf-sup constant degrades exactly as m^{-1} up to a positive multiplicative constant. This characterization shows that the failure of classical decoupled solvers in densely packed micro-channels is an intrinsic property of the PDE, not an implementation artifact, and exposes the a priori error amplification that standard finite element discretizations inherit.

Our sharp estimate also provides a quantitative foundation for stabilization. Advanced monolithic architectures—such as Augmented Lagrangian (AL) methods—bypass this geometric penalty entirely [18, 19, 20, 21] when the penalty parameter is scaled proportional to m^2 , as dictated by our pre-asymptotic bounds.

Our contributions are:

1. **Sharp bound.** We prove $\beta(\Omega_m) = \Theta(m^{-1})$ for the continuous LBB constant on generalized periodic pillar arrays (Theorem 1), the first explicit pre-asymptotic characterization for multiply-connected microfluidic geometries.
2. **Discretization consequence.** We map the sharp bound to its a priori error implication: the standard Brezzi estimate is amplified by $\beta_h^{-1} = \mathcal{O}(m)$, and we numerically expose an asymmetric pollution under macroscopic boundary conditions. As the mesh size h is refined such that $mh \sim \mathcal{O}(1)$, under Dirichlet inflow profiles the relative pressure error stagnates at $\mathcal{O}(1)$, whereas under Neumann pressure-drop conditions first-order convergence is retained.
3. **Algorithmic consequence.** We show that the pressure Schur complement satisfies $\kappa(M_p^{-1}S) \geq \mathcal{O}(m^2)$ and derive that setting the Augmented Lagrangian penalty as $\gamma \propto m^2$ restores $\mathcal{O}(1)$ outer iteration counts. The resulting parameter-free scaling reduces outer FGMRES iterations from 437 to 22 on a 1.85M-DoF square array and from 687 to 24 on a 1.77M-DoF DLD array.

Section 2 states the geometric setting and the main theorem. Section 3 establishes three lemmas. Section 4 proves the main theorem. Section 5 validates the $\Theta(m^{-1})$ rate, exposes its impact on discretization accuracy, and demonstrates density-independent AL convergence. Section 6 concludes.

2. Mathematical Model and Main Results

2.1. Geometric Configuration of tight domains with Periodic Pillars

Let the global bounding fluid domain, representing the macroscopic scale of the system, be $\Omega \subset \mathbb{R}^d$ ($d \in \{2, 3\}$), defined by a rectangular base cross-section $\Omega_{xy} = (0, L_x) \times (0, L_y)$. For $d = 2$, the domain is strictly planar ($\Omega = \Omega_{xy}$), while for $d = 3$, $\Omega = \Omega_{xy} \times (0, L_z)$ represents a microfluidic channel of height L_z . We assume that the domain is dimensionless, which means the lengths of L_x and L_y are of magnitude $\mathcal{O}(1)$.

To model the microscopic geometry of generic microfluidic pillar arrays, including DLD devices, we introduce a characteristic pore-scale periodicity parameter $\epsilon > 0$ and define a dimensionless reference

lattice in the x - y plane spanned by basis vectors $\hat{\mathbf{e}}_1 = (1, \delta)$ and $\hat{\mathbf{e}}_2 = (0, 1)$, where $\delta \in [0, 1)$ is the row shift fraction.

The physical lattice nodes are distributed in the cross-sectional plane according to:

$$\mathbf{x}_{i,j} = \epsilon(i\hat{\mathbf{e}}_1 + j\hat{\mathbf{e}}_2), \quad (i, j) \in \mathbb{Z}^2. \quad (1)$$

Let $P_0 \subset \mathbb{R}^2$ be the open, dimensionless reference cell in the cross-sectional plane centered at the origin, spanned by $\hat{\mathbf{e}}_1$ and $\hat{\mathbf{e}}_2$. We define a dimensionless reference pillar cross-section $T \subset \mathbb{R}^2$ as an arbitrary connected, closed, and bounded set with a Lipschitz boundary. To ensure that the pore-scale pillar cross-section is non-degenerate and guarantees a uniform local fluid neighborhood for subsequent theoretical restrictions, we assume there exist two concentric open balls $B(0, r_0)$ and $B(0, R_0)$ with geometric constants $0 < r_0 \leq R_0$ such that:

$$B(0, r_0) \subset T \subset B(0, R_0) \subset P_0. \quad (2)$$

Let $a_\epsilon > 0$ be a dimensional scaling parameter representing the characteristic size of the physical pillars. To prevent the pillars from overlapping the cell boundaries and to strictly bound the maximum solid volume fraction, we assume there exists a positive geometric constant $\kappa \in (0, 1)$ such that the relative pillar size uniformly satisfies $\sup_\epsilon(a_\epsilon/\epsilon) \leq \kappa$. The individual solid pillar at cell (i, j) is generated by scaling the reference pillar cross-section and translating it to the lattice node. Depending on the dimension d , it is either a 2D inclusion in the plane or a cylinder of cross-section $a_\epsilon T$ spanning the full channel height:

$$T_\epsilon^{i,j} = \begin{cases} \mathbf{x}_{i,j} + a_\epsilon T & \text{for } d = 2, \\ (\mathbf{x}_{i,j} + a_\epsilon T) \times [0, L_z] & \text{for } d = 3. \end{cases} \quad (3)$$

To rigorously avoid the truncation of pillars at the global lateral boundaries $\partial\Omega_{xy}$, we restrict the placement of solid pillars to the interior of the channel. We define the active index set $\mathcal{I}_\epsilon$ consisting of indices for which the corresponding projected physical cell $P_\epsilon^{i,j} = \mathbf{x}_{i,j} + \epsilon P_0$ is strictly contained within Ω_{xy} :

$$\mathcal{I}_\epsilon := \{(i, j) \in \mathbb{Z}^2 : \bar{P}_{i,j} \subset \Omega_{xy}\}. \quad (4)$$

Finally, the effective fluid domain, denoted by Ω_m (with $m \propto \epsilon^{-1}$ representing the pillar density in a unit length), is defined as the global bounding channel excluding the union of all fully contained solid pillars:

$$\Omega_m := \Omega \setminus \bigcup_{(i,j) \in \mathcal{I}_\epsilon} T_\epsilon^{i,j}. \quad (5)$$

This geometric arrangement is illustrated schematically in Figure 1, which shows the periodic parallelogram cells and the relative positions of the pillars within the array.




Figure 1: Geometric configuration of the periodic pillar array.

2.2. The Stokes Problem and the inf-sup Condition

We consider the steady, incompressible Stokes equations modeling the creeping flow of a viscous fluid through the effective fluid domain Ω_m :

$$\begin{cases} -\mu\Delta \mathbf{u} + \nabla p = f & \text{in } \Omega_m, \\ \nabla \cdot \mathbf{u} = 0 & \text{in } \Omega_m, \\ \mathbf{u} = 0 & \text{on } \partial\Omega_m, \end{cases} \quad (6)$$

where $\mathbf{u}$ is the fluid velocity, p is the pressure, $\mu > 0$ is the dynamic viscosity, and f represents the body force. The boundary $\partial\Omega_m = \partial\Omega \cup (\cup_{i,j} \partial T_\epsilon^{i,j})$ incorporates the no-slip condition on both the exterior channel walls and the surfaces of all internal pillars.

To cast (6) into a variational form, we introduce the standard functional spaces. The velocity space is $V(\Omega_m) = H_0^1(\Omega_m)^d$, equipped with the norm $\|\mathbf{u}\|_V = \|\nabla \mathbf{u}\|_{L^2(\Omega_m)}$. The pressure space is $Q(\Omega_m) = L_0^2(\Omega_m) = \{q \in L^2(\Omega_m) : \int_{\Omega_m} q \, dx = 0\}$, equipped with the standard L^2 -norm.

The well-posedness of the Stokes saddle-point problem, and the robust coupling between the velocity and pressure spaces, are governed by the classical Ladyzhenskaya-Babuška-Brezzi (LBB) condition. The continuous LBB constant for the perforated pillar-array domain Ω_m is defined as:

$$\beta(\Omega_m) := \inf_{q \in Q(\Omega_m) \setminus \{0\}} \sup_{\mathbf{v} \in V(\Omega_m) \setminus \{0\}} \frac{\int_{\Omega_m} q \nabla \cdot \mathbf{v} \, dx}{\|q\|_{L^2(\Omega_m)} \|\nabla \mathbf{v}\|_{L^2(\Omega_m)}}. \quad (7)$$

2.3. Main Results

The primary objective of this paper is to establish the sharp pre-asymptotic behavior of the continuous LBB constant $\beta(\Omega_m)$ as the pillar density increases (i.e., $\epsilon \rightarrow 0$ or equivalently $m \rightarrow \infty$). We characterize this severe degradation through the following theorem, the proofs of which are deferred to Section 4. Here and throughout the paper, the notation $\Theta(\cdot)$ denotes the two-sided asymptotic bound: for positive functions $f(x), g(x)$, we write $f(x) = \Theta(g(x))$ if there exist constants $c, C > 0$ and x_0 such that $c g(x) \leq f(x) \leq C g(x)$ for all $x \geq x_0$ (or as $x \rightarrow 0$), i.e., f grows or decays exactly at the rate of g up to positive multiplicative constants.

Theorem 1 (Sharp Pre-asymptotic Estimate of the Continuous LBB Constant). *Under the geometric assumptions defined in Section 2.1, assume the obstacle size scales proportionally with the periodicity (i.e., $a_\epsilon \propto \epsilon$). Then, the continuous inf-sup constant $\beta(\Omega_m)$ satisfies the following estimate:*

$$\beta(\Omega_m) = \Theta(\epsilon) = \Theta(m^{-1}). \quad (8)$$

3. Auxiliary Lemmas for the Stability Analysis

Before proving the main theorem, we establish three useful lemmas. From this section, we refer the term pillar as obstacle. Lemma 1 provides a global Poincaré inequality that explicitly captures the geometric penalty σ_ϵ induced by the obstacle array. Lemma 2 ensures the existence of a test pressure field with a uniformly bounded gradient-to-norm ratio, independent of the obstacle density. Finally, Lemma 3 constructs a restriction operator that enforces the microscopic no-slip conditions on global velocity fields while preserving their macroscopic divergence, providing the crucial stability estimates needed for the inf-sup lower bound. We give proof for 2-dimensional case, and the 3-dimensional case can be proved similarly.

Definition 1 (Asymptotic Parameter for Obstacle Arrays). *For the geometric setting in Section 2, the asymptotic behavior of the fluid flow is governed by the critical geometric parameter σ_ϵ , defined as:*

$$\sigma_\epsilon := \epsilon \left| \log \frac{a_\epsilon}{\epsilon} \right|^{\frac{1}{2}} \quad (9)$$

The following lemma generalizes the classical results from homogenization theory [16, 17] to our specific rectangular Bravais lattice configuration.

Lemma 1 (Global Poincaré Inequality). *Let Ω_m be the fluid domain with dense obstacles defined in Section 2, with zero boundary conditions on the holes $\partial T_\epsilon^{i,j}$ and the exterior boundary $\partial\Omega$. There exists a constant $C > 0$, independent of ϵ (and thus independent of the obstacle density), such that for all $\mathbf{v} \in H_0^1(\Omega_m)^d$:*

$$\|\mathbf{v}\|_{L^2(\Omega_m)} \leq C \min\{1, \sigma_\epsilon\} \|\nabla \mathbf{v}\|_{L^2(\Omega_m)}. \quad (10)$$

Proof. We extend $\mathbf{v} \in H_0^1(\Omega_m)^d$ by zero inside the obstacles $T_\epsilon^{i,j}$ to obtain $\tilde{\mathbf{v}} \in H_0^1(\Omega)^d$. The classical Poincaré inequality on the macroscopic domain Ω immediately yields $\|\tilde{\mathbf{v}}\|_{L^2(\Omega)} \leq C_P \|\nabla \tilde{\mathbf{v}}\|_{L^2(\Omega)}$, providing the global upper bound C_P independent of ϵ .

To derive the σ_ϵ -dependent local bound, we decompose Ω into the union of active fluid cells $P_\epsilon^{i,j}$ and a boundary layer Λ_ϵ near $\partial\Omega$ containing no complete obstacles. The width of Λ_ϵ is of order $\mathcal{O}(\epsilon)$. Since $\mathbf{v} = 0$ on $\partial\Omega$, the standard one-dimensional Poincaré inequality across the thickness of Λ_ϵ gives

$$\|\mathbf{v}\|_{L^2(\Lambda_\epsilon)}^2 \leq C_0 \epsilon^2 \|\nabla \mathbf{v}\|_{L^2(\Lambda_\epsilon)}^2. \quad (11)$$

For each cell $P_\epsilon^{i,j}$, let $B_{\text{out}}^{i,j}$ be its circumscribed ball with radius $R_\epsilon = C_{\text{out}}\epsilon$, and $B_{\text{in}}^{i,j} \subset T_\epsilon^{i,j}$ be the inscribed ball with radius $r_\epsilon = \alpha a_\epsilon$ (where $\alpha, C_{\text{out}} > 0$ are geometric lattice constants). Since $\mathbf{v} = 0$ on $\partial T_\epsilon^{i,j}$, it vanishes on $\partial B_{\text{in}}^{i,j}$. Using polar coordinates (r, θ) centered at the obstacle, the fundamental theorem of calculus along radial segments yields

$$\mathbf{v}(r, \theta) = \int_{r_\epsilon}^r \frac{\partial \mathbf{v}}{\partial \rho}(\rho, \theta) d\rho. \quad (12)$$

Applying the Cauchy-Schwarz inequality, we separate the derivative and the geometric weight:

$$|\mathbf{v}(r, \theta)|^2 \leq \left(\int_{r_\epsilon}^r \left| \frac{\partial \mathbf{v}}{\partial \rho} \right|^2 \rho d\rho \right) \left(\int_{r_\epsilon}^r \frac{1}{\rho} d\rho \right) \leq \ln \left(\frac{r}{r_\epsilon} \right) \int_{r_\epsilon}^{R_\epsilon} |\nabla \mathbf{v}|^2 \rho d\rho. \quad (13)$$

Integrating this point-wise bound over the outer ball $B_{\text{out}}^{i,j}$, we obtain

$$\begin{aligned} \|\mathbf{v}\|_{L^2(B_{\text{out}}^{i,j})}^2 &= \int_0^{2\pi} \int_{r_\epsilon}^{R_\epsilon} |\mathbf{v}(r, \theta)|^2 r dr d\theta \leq \|\nabla \mathbf{v}\|_{L^2(B_{\text{out}}^{i,j})}^2 \int_{r_\epsilon}^{R_\epsilon} r \ln \left(\frac{r}{r_\epsilon} \right) dr \\ &= \|\nabla \mathbf{v}\|_{L^2(B_{\text{out}}^{i,j})}^2 \left[\frac{R_\epsilon^2}{2} \ln \left(\frac{R_\epsilon}{r_\epsilon} \right) - \frac{R_\epsilon^2 - r_\epsilon^2}{4} \right]. \end{aligned} \quad (14)$$

Dropping the negative term and substituting the radii $R_\epsilon = C_{\text{out}}\epsilon$ and $r_\epsilon = \alpha a_\epsilon$, the geometric bracket is bounded by $\frac{1}{2} C_{\text{out}}^2 \epsilon^2 \ln \left(\frac{C_{\text{out}}\epsilon}{\alpha a_\epsilon} \right)$. By defining the characteristic scaling $\sigma_\epsilon^2 := \epsilon^2 |\ln(a_\epsilon/\epsilon)|$, there exists a constant $C_1 > 0$ such that the integral evaluates to $\leq C_1 \sigma_\epsilon^2$.

Summing over all cells $P_\epsilon^{i,j}$ and accounting for the finite overlap of the circumscribed balls, we have $\sum_{i,j} \|\mathbf{v}\|_{L^2(P_\epsilon^{i,j})}^2 \leq C_2 \sigma_\epsilon^2 \|\nabla \mathbf{v}\|_{L^2(\Omega_m)}^2$. Combining this with the boundary layer bound, we conclude

$$\|\mathbf{v}\|_{L^2(\Omega_m)}^2 = \|\mathbf{v}\|_{L^2(\Lambda_\epsilon)}^2 + \sum_{i,j} \|\mathbf{v}\|_{L^2(P_\epsilon^{i,j})}^2 \leq C_3 \sigma_\epsilon^2 \|\nabla \mathbf{v}\|_{L^2(\Omega_m)}^2. \quad (15)$$

Taking the square root and combining this microscopic limit with the macroscopic bound $\|\mathbf{v}\|_{L^2(\Omega_m)} \leq C_P \|\nabla \mathbf{v}\|_{L^2(\Omega_m)}$ completes the proof. $\square$

Lemma 2 (Existence of a Uniformly Bounded Test Pressure). *For the effective fluid domain Ω_m defined in Section 2, there exists a non-trivial test pressure field $q^* \in Q(\Omega_m) \cap H^1(\Omega_m)$ and a strictly positive constant $C_q > 0$, independent of the obstacle density m (and thus independent of ϵ), such that its gradient is bounded by its norm:*

$$\|\nabla q^*\|_{L^2(\Omega_m)} \leq C_q \|q^*\|_{L^2(\Omega_m)}. \quad (16)$$

Proof. Let $(x, y) \in \Omega = (0, L_x) \times (0, L_y)$ denote the spatial coordinates. Define the test pressure field $q^*(x, y) = x - \bar{x}$, where $\bar{x} = |\Omega_m|^{-1} \int_{\Omega_m} x dx dy$ is the spatial mean, ensuring $q^* \in Q(\Omega_m)$ (i.e., zero mean on the fluid domain). Since $\nabla q^* = (1, 0)^\top$, its gradient norm is exactly the measure of the fluid domain:

$$\|\nabla q^*\|_{L^2(\Omega_m)}^2 = \int_{\Omega_m} 1 dx dy = |\Omega_m|. \quad (17)$$

To bound $\|q^*\|_{L^2(\Omega_m)}$ from below, we evaluate the geometric variance of the x -coordinate over the perforated domain Ω_m :

$$\|q^*\|_{L^2(\Omega_m)}^2 = \int_{\Omega_m} (x - \bar{x})^2 dx dy. \quad (18)$$

By the Bathtub principle (or symmetric decreasing rearrangement), for any domain $\Omega_m \subset \Omega$ with a fixed total area $|\Omega_m|$, the integral in (18) is strictly minimized when Ω_m is rearranged into a single continuous

vertical strip centered at $\bar{x}$ with width $w = |\Omega_m|/L_y$. Thus, we obtain the rigorous geometric lower bound:

$$\int_{\Omega_m} (x - \bar{x})^2 dx dy \geq \int_0^{L_y} \int_{-w/2}^{w/2} \xi^2 d\xi dy = L_y \frac{w^3}{12} = \frac{|\Omega_m|^3}{12L_y^2}. \quad (19)$$

Taking the ratio of the gradient norm squared to the L^2 -norm squared, we immediately obtain a bound dependent purely on the fluid domain measure:

$$\frac{\|\nabla q^*\|_{L^2(\Omega_m)}^2}{\|q^*\|_{L^2(\Omega_m)}^2} \leq \frac{|\Omega_m|}{|\Omega_m|^3/(12L_y^2)} = \frac{12L_y^2}{|\Omega_m|^2}. \quad (20)$$

To ensure this bound is uniformly independent of the obstacle density, we assume the maximum solid area fraction is strictly bounded away from complete blockage, denoted by $\phi_{max} \in (0, 1)$. Indeed, ϕ_{max} is a fixed geometric constant strictly less than unity, as $T \subset P_0$ and $a_\epsilon/\epsilon \leq \kappa$. Physically, this ensures the existence of interconnected fluid channels and prevents the complete blockage of the device. Consequently, the effective fluid area satisfies $|\Omega_m| \geq (1 - \phi_{max})|\Omega| = (1 - \phi_{max})L_x L_y$. Substituting this macroscopic bound yields:

$$\frac{\|\nabla q^*\|_{L^2(\Omega_m)}}{\|q^*\|_{L^2(\Omega_m)}} \leq \frac{\sqrt{12}L_y}{|\Omega_m|} \leq \frac{\sqrt{12}}{(1 - \phi_{max})L_x} := C_q. \quad (21)$$

This completes the proof. $\square$

Lemma 3 (Properties of the Restriction Operator). *For the effective fluid domain Ω_m defined in Section 2, there exists a linear restriction operator $R_\epsilon : H_0^1(\Omega)^d \rightarrow H_0^1(\Omega_m)^d$ satisfying the following fundamental properties:*

1. *Preservation of zero-trace functions: For any $\mathbf{u} \in H_0^1(\Omega_m)^d$, extended by zero to Ω as $\tilde{\mathbf{u}}$, it holds that $R_\epsilon(\tilde{\mathbf{u}}) = \mathbf{u}$ in Ω_m .*
2. *Controlled modification of divergence: For any $\mathbf{u} \in H_0^1(\Omega)^d$, the divergence of the restricted field in Ω_m satisfies $\nabla \cdot R_\epsilon(\mathbf{u}) = \nabla \cdot \mathbf{u} + c_{i,j}$ within a local control volume $C_\epsilon^{i,j}$ surrounding each solid obstacle $T_\epsilon^{i,j}$, and $\nabla \cdot R_\epsilon(\mathbf{u}) = \nabla \cdot \mathbf{u}$ elsewhere. Crucially, if $\nabla \cdot \mathbf{u} = 0$ identically inside the solid obstacles, then $\nabla \cdot R_\epsilon(\mathbf{u}) = \nabla \cdot \mathbf{u}$ everywhere in Ω_m .*
3. *Stability estimate: For any $\mathbf{u} \in H_0^1(\Omega)^d$, the gradient of the restricted field satisfies the bound:*

$$\|\nabla R_\epsilon(\mathbf{u})\|_{L^2(\Omega_m)} \leq C (\|\nabla \mathbf{u}\|_{L^2(\Omega)} + (1 + \sigma_\epsilon^{-1})\|\mathbf{u}\|_{L^2(\Omega)}) \quad (22)$$

where $C > 0$ is a constant independent of ϵ .

Proof of Lemma 3. We construct the restriction operator R_ϵ by solving local Stokes problems. For each active cell $(i, j) \in \mathcal{I}_\epsilon$, let $B_\epsilon^{i,j} \subset P_\epsilon^{i,j}$ be an inscribed open ball strictly covering the solid obstacle $T_\epsilon^{i,j}$, and define the fluid control volume $C_\epsilon^{i,j} = B_\epsilon^{i,j} \setminus T_\epsilon^{i,j}$.

For any $\mathbf{u} \in H_0^1(\Omega)^d$, we set $R_\epsilon(\mathbf{u})$ piecewise. Away from the obstacles:

$$R_\epsilon(\mathbf{u}) = \mathbf{u} \quad \text{in } \Omega \setminus \bigcup_{(i,j) \in \mathcal{I}_\epsilon} B_\epsilon^{i,j}, \quad (23)$$

and inside the solid obstacles:

$$R_\epsilon(\mathbf{u}) = 0 \quad \text{in } T_\epsilon^{i,j}. \quad (24)$$

Inside $C_\epsilon^{i,j}$, we define $R_\epsilon(\mathbf{u}) = \mathbf{v}$ as the solution to the local Stokes problem:

$$\begin{cases} -\Delta \mathbf{v} + \nabla q = -\Delta \mathbf{u} & \text{in } C_\epsilon^{i,j} \\ \nabla \cdot \mathbf{v} = \nabla \cdot \mathbf{u} + c_{i,j} & \text{in } C_\epsilon^{i,j} \\ \mathbf{v} = \mathbf{u} & \text{on } \partial B_\epsilon^{i,j} \\ \mathbf{v} = 0 & \text{on } \partial T_\epsilon^{i,j} \end{cases} \quad (25)$$

where $c_{i,j}$ is a constant. Well-posedness requires the compatibility condition:

$$\int_{C_\epsilon^{i,j}} (\nabla \cdot \mathbf{v}) dx = \int_{\partial C_\epsilon^{i,j}} \mathbf{v} \cdot \mathbf{n} ds = \int_{\partial B_\epsilon^{i,j}} \mathbf{u} \cdot \mathbf{n}_{\text{out}} ds - \int_{\partial T_\epsilon^{i,j}} 0 \cdot \mathbf{n}_{\text{in}} ds. \quad (26)$$

Applying the divergence theorem to $\mathbf{u}$ over $B_\epsilon^{i,j}$ yields:

$$\int_{\partial B_\epsilon^{i,j}} \mathbf{u} \cdot \mathbf{n}_{\text{out}} \, ds = \int_{B_\epsilon^{i,j}} \nabla \cdot \mathbf{u} \, dx = \int_{C_\epsilon^{i,j}} \nabla \cdot \mathbf{u} \, dx + \int_{T_\epsilon^{i,j}} \nabla \cdot \mathbf{u} \, dx. \quad (27)$$

Substituting (27) into (26) and using the divergence equation from (25), we obtain:

$$\int_{C_\epsilon^{i,j}} (\nabla \cdot \mathbf{u} + c_{i,j}) \, dx = \int_{C_\epsilon^{i,j}} \nabla \cdot \mathbf{u} \, dx + \int_{T_\epsilon^{i,j}} \nabla \cdot \mathbf{u} \, dx. \quad (28)$$

This uniquely determines $c_{i,j}$:

$$c_{i,j} = \frac{1}{|C_\epsilon^{i,j}|} \int_{T_\epsilon^{i,j}} \nabla \cdot \mathbf{u} \, dx, \quad (29)$$

which establishes Property 2.

For Property 1, if $\mathbf{u} \in H_0^1(\Omega_m)^d$ is extended by zero to Ω , then $\mathbf{u} = 0$ on $\partial T_\epsilon^{i,j}$ and $\nabla \cdot \mathbf{u} = 0$ identically in $T_\epsilon^{i,j}$. Consequently, $c_{i,j} = 0$, and (25) admits the unique solution $\mathbf{v} = \mathbf{u}$ in $C_\epsilon^{i,j}$.

For Property 3, we rescale the physical ball $B_\epsilon^{i,j}$ to a unit ball $\hat{B}$ via $y = (x - x_{i,j})/\epsilon$. The obstacle transforms into $\hat{T}_\eta$ with characteristic size $\eta = a_\epsilon/\epsilon$, defining the reference annulus $\hat{C}_\eta = \hat{B} \setminus \hat{T}_\eta$. Defining $\hat{\mathbf{u}}(y) = \mathbf{u}(x_{i,j} + \epsilon y)$ and $\hat{\mathbf{v}}(y) = \mathbf{v}(x_{i,j} + \epsilon y)$, standard Bogovskii estimates [16, Lemma 2.2.5] yield:

$$\|\hat{\nabla} \hat{\mathbf{v}}\|_{L^2(\hat{C}_\eta)}^2 \leq C \left(\|\hat{\nabla} \hat{\mathbf{u}}\|_{L^2(\hat{B})}^2 + |\log \eta|^{-1} \|\hat{\mathbf{u}}\|_{L^2(\hat{B})}^2 \right). \quad (30)$$

Scaling back to physical variables, we use $\|\hat{\nabla} \hat{\mathbf{v}}\|_{L^2(\hat{C}_\eta)}^2 = \|\nabla \mathbf{v}\|_{L^2(C_\epsilon^{i,j})}^2$, $\|\hat{\nabla} \hat{\mathbf{u}}\|_{L^2(\hat{B})}^2 = \|\nabla \mathbf{u}\|_{L^2(B_\epsilon^{i,j})}^2$, and $\|\hat{\mathbf{u}}\|_{L^2(\hat{B})}^2 = \epsilon^{-2} \|\mathbf{u}\|_{L^2(B_\epsilon^{i,j})}^2$:

$$\|\nabla \mathbf{v}\|_{L^2(C_\epsilon^{i,j})}^2 \leq C \left(\|\nabla \mathbf{u}\|_{L^2(B_\epsilon^{i,j})}^2 + \frac{|\log \eta|^{-1}}{\epsilon^2} \|\mathbf{u}\|_{L^2(B_\epsilon^{i,j})}^2 \right). \quad (31)$$

Recalling $\eta = a_\epsilon/\epsilon$ and $\sigma_\epsilon = \epsilon |\log(a_\epsilon/\epsilon)|^{1/2}$, the penalty coefficient scales as:

$$\frac{|\log \eta|^{-1}}{\epsilon^2} \propto \frac{|\log(a_\epsilon/\epsilon)|^{-1}}{\epsilon^2} = \sigma_\epsilon^{-2}, \quad (32)$$

yielding the local estimate:

$$\|\nabla \mathbf{v}\|_{L^2(C_\epsilon^{i,j})}^2 \leq C \left(\|\nabla \mathbf{u}\|_{L^2(B_\epsilon^{i,j})}^2 + \sigma_\epsilon^{-2} \|\mathbf{u}\|_{L^2(B_\epsilon^{i,j})}^2 \right). \quad (33)$$

Summing over all active cells $(i,j) \in \mathcal{I}_\epsilon$ and noting the topologically bounded finite overlap of $B_\epsilon^{i,j}$, we obtain the global bound:

$$\|\nabla R_\epsilon(\mathbf{u})\|_{L^2(\Omega_m)}^2 \leq C' \left(\|\nabla \mathbf{u}\|_{L^2(\Omega)}^2 + \sigma_\epsilon^{-2} \|\mathbf{u}\|_{L^2(\Omega)}^2 \right). \quad (34)$$

Taking the square root and applying $\sqrt{a^2 + b^2} \leq a + b$ yields the final stability estimate:

$$\|\nabla R_\epsilon(\mathbf{u})\|_{L^2(\Omega_m)} \leq C'' \left(\|\nabla \mathbf{u}\|_{L^2(\Omega)} + \sigma_\epsilon^{-1} \|\mathbf{u}\|_{L^2(\Omega)} \right), \quad (35)$$

completing the proof. $\square$

4. Proofs of the Main Theorems

With the foundational geometric and functional estimates established in Section 3, we now present the rigorous proofs for the sharp pre-asymptotic degradation rate of the pressure-velocity coupling stability in the obstacle array domain Ω_m .

4.1. Proof of Theorem 1: The Upper Bound

Proof of the upper bound. The continuous inf-sup constant for the domain Ω_m is defined by the variational problem:

$$\beta(\Omega_m) = \inf_{q \in L_0^2(\Omega_m) \setminus \{0\}} \sup_{\mathbf{v} \in H_0^1(\Omega_m)^d \setminus \{\mathbf{0}\}} \frac{\int_{\Omega_m} q \nabla \cdot \mathbf{v} \, dx}{\|q\|_{L^2(\Omega_m)} \|\nabla \mathbf{v}\|_{L^2(\Omega_m)}}. \quad (36)$$

To establish an upper bound, it suffices to evaluate the supremum quotient for a specific test pressure field. We select the global linear pressure field $q^*(x, y) = x - \bar{x} \in Q(\Omega_m)$ constructed in Lemma 2.

By integration by parts, we get the following estimate.

$$\int_{\Omega_m} q^* \nabla \cdot \mathbf{v} \, dx = - \int_{\Omega_m} \nabla q^* \cdot \mathbf{v} \, dx \leq \|\nabla q^*\|_{L^2(\Omega_m)} \|\mathbf{v}\|_{L^2(\Omega_m)} \leq C_P \sigma_\epsilon \|\nabla q^*\|_{L^2(\Omega_m)} \|\nabla \mathbf{v}\|_{L^2(\Omega_m)}, \quad (37)$$

where we used the global Poincaré inequality from Lemma 1. We insert this bound into the quotient in (7) evaluated at q^* :

$$\beta(\Omega_m) \leq \sup_{\mathbf{v} \in H_0^1(\Omega_m)^d} \frac{C_P \sigma_\epsilon \|\nabla q^*\|_{L^2(\Omega_m)} \|\nabla \mathbf{v}\|_{L^2(\Omega_m)}}{\|q^*\|_{L^2(\Omega_m)} \|\nabla \mathbf{v}\|_{L^2(\Omega_m)}} = C_P \sigma_\epsilon \frac{\|\nabla q^*\|_{L^2(\Omega_m)}}{\|q^*\|_{L^2(\Omega_m)}}. \quad (38)$$

By Lemma 2, the ratio $\|\nabla q^*\|_{L^2(\Omega_m)} / \|q^*\|_{L^2(\Omega_m)}$ is bounded by a constant C_q strictly independent of the obstacle density. Therefore:

$$\beta(\Omega_m) \leq (C_P C_q) \sigma_\epsilon. \quad (39)$$

For the specific case where the obstacle size scales proportionally with the cell size ($a_\epsilon = \kappa \epsilon$), the parameter simplifies to $\sigma_\epsilon = \epsilon |\log \kappa|^{1/2} = \mathcal{O}(\epsilon)$. Since the number of cells per dimension is $m \propto \epsilon^{-1}$, we conclude that $\beta(\Omega_m) \leq \mathcal{O}(m^{-1})$, completing the proof. $\square$

4.2. Proof of Theorem 1: The Lower Bound and Sharpness

Proof of lower bound. Proving a lower bound for the continuous inf-sup constant $\beta(\Omega_m) \geq C^*$ is equivalent to showing that, for any $q \in L_0^2(\Omega_m)$, there exists a velocity field $\mathbf{v} \in H_0^1(\Omega_m)^d$ such that $\nabla \cdot \mathbf{v} = q$ and $\|\nabla \mathbf{v}\|_{L^2(\Omega_m)} \leq (C^*)^{-1} \|q\|_{L^2(\Omega_m)}$.

Given $q \in L_0^2(\Omega_m)$, we extend it by zero to the entire macroscopic domain Ω to define $\tilde{q}$. Since the extension is by zero, $\tilde{q} \in L_0^2(\Omega)$ and $\|\tilde{q}\|_{L^2(\Omega)} = \|q\|_{L^2(\Omega_m)}$. Because the macroscopic domain Ω is fixed and independent of ϵ , standard estimate guarantees the existence of a global velocity field $\mathbf{V} \in H_0^1(\Omega)^d$ satisfying:

$$\nabla \cdot \mathbf{V} = \tilde{q} \quad \text{in } \Omega, \quad \text{and} \quad \|\nabla \mathbf{V}\|_{L^2(\Omega)} \leq C_B \|\tilde{q}\|_{L^2(\Omega)} = C_B \|q\|_{L^2(\Omega_m)}, \quad (40)$$

where C_B is the standard LBB constant for the unperforated domain Ω .

To enforce the no-slip condition on the obstacle boundaries while preserving the target divergence, we apply the restriction operator constructed in Lemma 3, defining $\mathbf{v} = R_\epsilon(\mathbf{V}) \in H_0^1(\Omega_m)^d$.

According to Property 2 of Lemma 3, because $\nabla \cdot \mathbf{V} = \tilde{q} = 0$ identically inside all solid obstacles $T_\epsilon^{i,j}$, the restriction operator perfectly preserves the divergence field globally without requiring local corrections (i.e., $c_{i,j} = 0$). Thus, $\nabla \cdot \mathbf{v} = \nabla \cdot \mathbf{V} = q$ everywhere in Ω_m .

To evaluate the norm of the constructed field, we invoke the stability estimate (Property 3) of Lemma 3:

$$\|\nabla \mathbf{v}\|_{L^2(\Omega_m)} \leq C_R (\|\nabla \mathbf{V}\|_{L^2(\Omega)} + (1 + \sigma_\epsilon^{-1}) \|\mathbf{V}\|_{L^2(\Omega)}). \quad (41)$$

We bound the L^2 -norm of the macroscopic field V using the standard Poincaré inequality on the fixed domain Ω , yielding $\|\mathbf{V}\|_{L^2(\Omega)} \leq C_0 \|\nabla \mathbf{V}\|_{L^2(\Omega)}$. Substituting this and the global infsup constant C_B into the stability estimate yields:

$$\begin{aligned} \|\nabla \mathbf{v}\|_{L^2(\Omega_m)} &\leq C_R (1 + C_0(1 + \sigma_\epsilon^{-1})) \|\nabla \mathbf{V}\|_{L^2(\Omega)} \\ &\leq C_R C_B (1 + C_0 + C_0 \sigma_\epsilon^{-1}) \|q\|_{L^2(\Omega_m)}. \end{aligned} \quad (42)$$

For a small ϵ (and consequently a large σ_ϵ^{-1}), the scaling term σ_ϵ^{-1} dominates the constant terms. Therefore, there exists an asymptotic constant $C' > 0$ such that:

$$\|\nabla \mathbf{v}\|_{L^2(\Omega_m)} \leq C' \sigma_\epsilon^{-1} \|q\|_{L^2(\Omega_m)}. \quad (43)$$

This establishes that the norm of the continuous Bogovskii operator on the perforated domain Ω_m is bounded from above by $\mathcal{O}(\sigma_\epsilon^{-1})$. Inverting this norm bound immediately provides the lower bound for the inf-sup constant:

$$\beta(\Omega_m) = \inf_q \sup_{\mathbf{v}} \frac{\int_{\Omega_m} q \nabla \cdot \mathbf{v}}{\|q\| \|\nabla \mathbf{v}\|} \geq \frac{1}{C' \sigma_\epsilon^{-1}} = C'' \sigma_\epsilon. \quad (44)$$

Combined with the upper bound $\beta(\Omega_m) \leq (C_P C_q) \sigma_\epsilon$ from Theorem 1, we conclude the sharp estimate $\beta(\Omega_m) = \Theta(\sigma_\epsilon)$. For $a_\epsilon \propto \epsilon$, this yields $\beta(\Omega_m) = \Theta(\epsilon) = \Theta(m^{-1})$, completing the proof. $\square$

Remark 1 (Discrete inf-sup Degradation). *The continuous degradation rate $\beta(\Omega_m) = \Theta(m^{-1})$ is directly inherited by the discrete saddle-point system. For stable conforming mixed finite elements (e.g., Taylor-Hood) admitting a mesh-independent Fortin operator, the discrete inf-sup constant is bounded below by $\beta_h \gtrsim \beta(\Omega_m) = \Theta(m^{-1})$ [6, 22]. Conversely, provided the discrete pressure space Q_h contains global linear polynomials, evaluating the discrete inf-sup quotient with $q^*(x, y)$ yields the sharp upper bound $\beta_h \leq \mathcal{O}(m^{-1})$. Thus, for standard stable discretizations, the discrete inf-sup constant degradation remains $\beta_h = \Theta(m^{-1})$. Detailed proofs are omitted as they follow straightforwardly from standard mixed finite element theory.*

5. Numerical Consequences of the Sharp inf-sup Bound

In this section, we present a series of numerical experiments designed to verify the sharp pre-asymptotic estimate of the continuous LBB constant established in Theorem 1 and to map this geometric penalty onto its two practical manifestations: the amplification of the a priori discretization error and the stagnation of standard iterative solvers. We then evaluate the robustness of the proposed adaptively scaled Augmented Lagrangian (AL) preconditioner as the pillar density m increases.

5.1. Discrete Formulation and Matrix Definitions

To strictly isolate the continuous, geometry-induced inf-sup degradation from any artificial numerical instabilities, we employ the standard Galerkin finite element method using the inf-sup stable Taylor-Hood element pair $(\mathbb{P}_2 - \mathbb{P}_1)$. This classical choice guarantees that the discrete LBB condition is inherently satisfied, ensuring that any spectral deterioration observed in the algebraic system originates purely from the micro-structural topology rather than spatial discretization artifacts.

Let $\mathcal{T}_h$ be a shape-regular finite element triangulation of the domain. We define the conforming discrete velocity and pressure spaces as:

$$V_h = \{\mathbf{v}_h \in C^0(\overline{\Omega})^d : \mathbf{v}_h|_K \in [\mathbb{P}_2(K)]^d, \forall K \in \mathcal{T}_h\} \cap W_0^{1,2}(\Omega)^d, \quad (45)$$

$$Q_h = \{q_h \in C^0(\overline{\Omega}) : q_h|_K \in \mathbb{P}_1(K), \forall K \in \mathcal{T}_h\} \cap L_0^2(\Omega), \quad (46)$$

where $d \in \{2, 3\}$ is the spatial dimension.

The discrete variational formulation seeks a velocity-pressure pair $(\mathbf{u}_h, p_h) \in V_h \times Q_h$ such that:

$$a(\mathbf{u}_h, \mathbf{v}_h) + b(\mathbf{v}_h, p_h) = (f, \mathbf{v}_h), \quad \forall \mathbf{v}_h \in V_h, \quad (47)$$

$$b(\mathbf{u}_h, q_h) = 0, \quad \forall q_h \in Q_h. \quad (48)$$

Here, $b(\mathbf{v}, q) = -\int q \nabla \cdot \mathbf{v} dx$ represents the divergence constraint, and the bilinear form $a(\mathbf{u}, \mathbf{v}) = \int \mu \nabla \mathbf{u} : \nabla \mathbf{v} dx$ encapsulates the viscous dissipation.

Upon choosing standard nodal basis functions for the spaces V_h and Q_h , this formulation leads to the classical generalized saddle-point algebraic system:

$$\mathcal{A} \mathbf{x} = \mathbf{b} \quad \implies \quad \begin{pmatrix} A & B^T \\ B & 0 \end{pmatrix} \begin{pmatrix} \mathbf{u} \\ \mathbf{p} \end{pmatrix} = \begin{pmatrix} \mathbf{f} \\ \mathbf{0} \end{pmatrix}, \quad (49)$$

where $\mathbf{u} \in \mathbb{R}^{n_u}$ and $\mathbf{p} \in \mathbb{R}^{n_p}$ denote the vectors of velocity and pressure degrees of freedom, respectively. The global block matrices are defined as follows:

- $A \in \mathbb{R}^{n_u \times n_u}$: the symmetric positive-definite velocity stiffness matrix representing the discrete operator associated with $a(\mathbf{u}, \mathbf{v})$. In the context of our perforated domain analysis, the condition number of this block inherently depends on the obstacle density.
- $B \in \mathbb{R}^{n_p \times n_u}$: the discrete negative divergence matrix associated with $b(\mathbf{v}, q)$.

- $B^T \in \mathbb{R}^{n_u \times n_p}$: the discrete gradient matrix.

With this algebraic framework established, the pressure Schur complement, explicitly defined as $S = BA^{-1}B^T$, becomes the central object of our pre-asymptotic stability analysis.

The primary objective of this section is twofold: first, to numerically validate the exact $\Theta(m^{-1})$ pre-asymptotic degradation rate of the discrete inf-sup constant established in our theoretical analysis, which translates algebraically to the minimal generalized eigenvalue of the scaled Schur complement; second, to inspect the behavior of the discrete solution as the pillar density increases, which is dominated not only by the quasi-optimal convergence, but also by the inf-sup constant; and third, to demonstrate how this analytical bound provides a rigorous foundation for choosing the penalty parameter in Augmented Lagrangian (AL) stabilization.

5.2. Discrete inf-sup Constant Degrades as $\Theta(m^{-1})$

To verify Theorem 1, we compute the discrete inf-sup constant β_h for a sequence of domains Ω_m with increasing pillar density m . The discrete constant is evaluated by solving the generalized eigenvalue problem $BA^{-1}B^Tq = \lambda M_p q$, where M_p is the pressure mass matrix, and extracting the square root of the smallest non-zero eigenvalue: $\beta_h = \sqrt{\lambda_{\min \neq 0}}$.

We consider a unit global bounding channel $\Omega = (0, 1)^2$ and partition it into an $m \times m$ unshifted Bravais lattice ($\delta = 0$). Circular pillars with a fixed packing ratio $\kappa = 0.25$ are placed at the center of each active cell. We utilize standard Taylor-Hood ($P_2 - P_1$) finite elements, ensuring that the mesh size h is refined sufficiently at each density level to capture the local shear layers.




Figure 2: Log-log plot of the discrete inf-sup constant β_h as a function of the pillar density m . The computed values precisely track the theoretical $\Theta(m^{-1})$ asymptote (indicated by the reference line with slope -1), confirming the sharpness of the pre-asymptotic degradation.

As illustrated in Figure 2, the computed discrete constants align with the theoretical $\Theta(m^{-1})$ trajectory. This numerical evidence corroborates our proof that the deterioration of the saddle-point stability is governed by the increasing pillar density, and that this geometric penalty is captured by the discrete inf-sup constant.

The discrete inf-sup constant tracks the theoretical $\Theta(m^{-1})$ slope across all tested densities. The geometric penalty is fully inherited by the discrete system; it is not an artifact of the discretization.

5.3. inf-sup Degradation Amplifies Pressure Error Under Dirichlet Constraints

5.3.1. Experimental Setup and Baseline Convergence

To investigate the implications of the stability degradation $\beta_h = \Theta(m^{-1})$ established in our theoretical analysis, we design two strictly controlled physical experiments. We consider Stokes flow through a deterministic lateral displacement (DLD) micro-pillar array without volumetric body forces ($\mathbf{f} = \mathbf{0}$). Let Γ_{in} , Γ_{out} , and Γ_{wall} denote the inlet, outlet, and no-slip boundary (including channel walls and micro-pillars), respectively. The computational domain is defined in Figure 3. To comprehensively test the system, we investigate two distinct macroscopic driving mechanisms:




Figure 3: Schematic of the experimental setup for the DLD micro-pillar array. The computational domain Ω is defined as $(0, L_x) \times (0, L_y)$, where $L_x = 2.0$ and $L_y = 1.0$. The array consists of $2m \times m$ unit cells, each containing a circular pillar of radius $r = 0.3\epsilon$ located at $(\epsilon(i + 0.5), \epsilon(j + 0.5))$ for $i = 0, 1, \dots, 2m - 1$ and $j = 0, 1, \dots, m - 1$, with $\epsilon = 1/m$. The interior evaluation subdomain Ω_{inner} is defined as $(0.25L_x, 0.75L_x) \times (0, L_y)$.

1. **Setup A: Pressure-Driven Flow.** We impose a macroscopic pressure drop ($\Delta p = m^2$) via normal stress Neumann boundary conditions at the inlet and outlet:

$$\begin{aligned} (-p\mathbf{I} + \nabla \mathbf{u}) \cdot \mathbf{n} &= -p_{\text{in}} \mathbf{n} \quad \text{on } \Gamma_{\text{in}}, \\ (-p\mathbf{I} + \nabla \mathbf{u}) \cdot \mathbf{n} &= -p_{\text{out}} \mathbf{n} \quad \text{on } \Gamma_{\text{out}}, \end{aligned} \quad (50)$$

where $p_{\text{in}} = m^2$, $p_{\text{out}} = 0.0$, and $\mathbf{n}$ is the outward unit normal vector.

2. **Setup B: Velocity-Driven Pipe Flow.** We impose a fixed macroscopic parabolic velocity profile at both the inlet and the outlet (pure Dirichlet conditions) to enforce mass conservation:

$$\mathbf{u} = \mathbf{u}_D(y) \quad \text{on } \Gamma_{\text{in}} \cup \Gamma_{\text{out}}, \quad (51)$$

where $\mathbf{u}_D(y) = [4U_{\text{max}}y(L_y - y)/L_y^2, 0]^T$ with $U_{\text{max}} = 1.0$.

In both cases, standard no-slip Dirichlet conditions are strictly applied to all physical obstacles and bounding walls:

$$\mathbf{u} = \mathbf{0} \quad \text{on } \Gamma_{\text{wall}}. \quad (52)$$

To isolate the discretization error induced by the degraded inf-sup stability from geometric approximation errors (i.e., variational crimes [23]), we adopt a hierarchical polynomial elevation strategy [24]. The reference true solution $(\mathbf{u}_{\text{ref}}, p_{\text{ref}})$ is computed using the high-order $P_3 - P_2$ element on the exactly identical second-order isoparametric curved mesh as the targeted Taylor-Hood $(P_2 - P_1)$ approximation $(\mathbf{u}_h, p_h)$. Relative and absolute errors are strictly evaluated within an interior macroscopic subdomain $\Omega_{\text{inner}} := (0.25L_x, 0.75L_x) \times (0, L_y)$ to bypass corner singularities.

Specifically, the relative errors for the velocity L^2 -norm, velocity H^1 -seminorm, and pressure L^2 -norm are defined respectively as:

$$e_{\mathbf{u}, L^2} = \frac{\|\mathbf{u}_h - \mathbf{u}_{\text{ref}}\|_{L^2(\Omega_{\text{inner}})}}{\|\mathbf{u}_{\text{ref}}\|_{L^2(\Omega_{\text{inner}})}}, \quad e_{\mathbf{u}, H^1} = \frac{\|\nabla \mathbf{u}_h - \nabla \mathbf{u}_{\text{ref}}\|_{L^2(\Omega_{\text{inner}})}}{\|\nabla \mathbf{u}_{\text{ref}}\|_{L^2(\Omega_{\text{inner}})}}, \quad (53)$$

$$e_{p, L^2} = \frac{\|(p_h - \bar{p}_h) - (p_{\text{ref}} - \bar{p}_{\text{ref}})\|_{L^2(\Omega_{\text{inner}})}}{\|p_{\text{ref}} - \bar{p}_{\text{ref}}\|_{L^2(\Omega_{\text{inner}})}}, \quad (54)$$

where $\bar{p} = \frac{1}{|\Omega_{\text{inner}}|} \int_{\Omega_{\text{inner}}} p \, dx$ denotes the spatial mean of the pressure over the evaluation subdomain. This zero-mean adjustment for the pressure field ensures a rigorous comparison by explicitly filtering out the arbitrary constant associated with the pure Dirichlet boundary conditions in Setup B.

To systematically control the mesh resolution across varying obstacle densities, we introduce a dimensionless grid factor, N_g . This factor dictates the local resolution within each microscopic unit cell, enforcing the global maximum mesh size to scale as $h_{\text{max}} \approx \epsilon/N_g = (m \cdot N_g)^{-1}$, where $\epsilon = 1/m$ is the characteristic pore size.

As demonstrated in Validation 1 (Table 1), for a fixed obstacle density ($m = 4$), successive mesh refinements yield the optimal asymptotic convergence rates regardless of the driving boundary conditions. The relative pressure error decays at $\mathcal{O}(h^2)$, while the velocity L^2 and H^1 errors decay at $\mathcal{O}(h^3)$ and $\mathcal{O}(h^2)$ respectively. This confirms the baseline consistency of the spatial discretization when the macroscopic structure is invariant.

Table 1: Validation 1: Baseline convergence with a fixed obstacle density ($m = 4$). Optimal Taylor-Hood convergence rates are achieved under both pressure-driven and velocity-driven boundary conditions as $h \rightarrow 0$.

Setup A: Pressure-Driven Boundary Conditions						
$h_{\max}$	Rel. p (L^2)	Order (p)	Rel. $\mathbf{u}$ (L^2)	Order ($\mathbf{u}_{L^2}$)	Rel. $\mathbf{u}$ (H^1)	Order ($\mathbf{u}_{H^1}$)
0.0786	1.90e-02	–	3.55e-02	–	1.70e-01	–
0.0549	7.85e-03	2.46	1.22e-02	2.98	9.81e-02	1.53
0.0426	4.71e-03	2.02	6.42e-03	2.53	6.56e-02	1.58
0.0338	3.25e-03	1.60	3.89e-03	2.16	4.75e-02	1.40
0.0284	2.17e-03	2.31	2.27e-03	3.10	3.40e-02	1.91
0.0247	1.55e-03	2.41	1.42e-03	3.35	2.52e-02	2.14
0.0218	1.18e-03	2.18	9.83e-04	2.94	1.99e-02	1.89
Setup B: Velocity-Driven Pipe Flow Boundary Conditions						
0.0786	1.90e-02	–	3.55e-02	–	1.70e-01	–
0.0549	8.48e-03	2.25	1.19e-02	3.05	9.83e-02	1.53
0.0426	5.24e-03	1.89	6.15e-03	2.59	6.57e-02	1.58
0.0338	3.50e-03	1.74	3.74e-03	2.15	4.75e-02	1.40
0.0284	2.29e-03	2.44	2.20e-03	3.06	3.40e-02	1.92
0.0247	1.60e-03	2.55	1.38e-03	3.34	2.52e-02	2.15
0.0218	1.21e-03	2.25	9.57e-04	2.91	1.99e-02	1.89

5.3.2. Amplification Penalty and Scaling Analysis

The pronounced degradation of the convergence emerges in Validation 2 (Tables 2 and 3). Here, we fix the local mesh resolution within each pore space ($N_g = 6$) and increase the obstacle density m . This strictly forces the global maximum mesh size to scale inversely with the density ($h_{\max} \propto m^{-1}$).

To prove that the observed loss of optimal convergence is not a numerical artifact but a fundamental mathematical manifestation of the multiscale saddle-point system, we analyze the error in three steps, mapping our theoretical bounds directly to the empirical data in Table 2.

Step 1: Absolute Scaling of the Physical Fields. Before analyzing the errors, we must establish the physical baselines driven by different boundary conditions. Crucially, we designed both setups to maintain equivalent macroscopic scalings, isolating the boundary formulation as the sole independent variable.

- **Setup A (Pressure-Driven):** By imposing a macroscopic pressure drop that explicitly scales with the characteristic hydraulic resistance ($\Delta p = m^2$), we artificially drive a macroscopically stable flow rate. According to Darcy’s law, the absolute velocity remains strictly $\mathcal{O}(1)$, while the reference pressure scales as $\mathcal{O}(m^2)$. This is verified in Table 2 (Setup A), where $\|\mathbf{u}_{\text{ref}}\|_{L^2}$ remains remarkably stable at $\approx 8.0 \times 10^{-3}$ across all densities up to $m = 128$, and $\|p_{\text{ref}}\|_{L^2}$ scales strictly quadratically, reaching 2.00×10^3 at $m = 128$.
- **Setup B (Velocity-Driven):** Forcing a constant macroscopic flow rate ($\|\mathbf{u}_{\text{ref}}\|_{L^2} \approx 1.0$) through densely packed pillars similarly reflects Darcy’s law. The required absolute pressure must identically grow quadratically: $\|p_{\text{ref}}\|_{L^2} = \mathcal{O}(m^2)$. Table 2 confirms this: as m doubles iteratively up to 128, $\|p_{\text{ref}}\|_{L^2}$ quadruples at each step, culminating at 2.44×10^5 .

Step 2: Invariance of the Relative Velocity Error. In both physical regimes, the spatial derivatives of the highly localized velocity field scale strictly with the microscopic length ($\epsilon \propto m^{-1}$):

$$|\mathbf{u}|_{H^k} \sim \epsilon^{-k} \|\mathbf{u}\|_{L^2} \sim m^k \|\mathbf{u}\|_{L^2}. \quad (55)$$

Because the local resolution within each pore space is fixed, the global mesh size satisfies $h \propto m^{-1}$. By standard interpolation theory for P_2 elements, the absolute best-approximation errors for the velocity field in the L^2 - and H^1 -norms are bounded respectively by:

$$\begin{aligned} \|\mathbf{u} - \mathbf{u}_h\|_{L^2} &\leq Ch^3 |\mathbf{u}|_{H^3} \sim (m^{-1})^3 (m^3 \|\mathbf{u}\|_{L^2}) \sim \|\mathbf{u}\|_{L^2}, \\ |\mathbf{u} - \mathbf{u}_h|_{H^1} &\leq Ch^2 |\mathbf{u}|_{H^3} \sim (m^{-1})^2 (m^3 \|\mathbf{u}\|_{L^2}) = m \|\mathbf{u}\|_{L^2} \sim |\mathbf{u}|_{H^1}. \end{aligned} \quad (56)$$

Dividing these absolute errors by their corresponding norms mathematically demonstrates that the relative velocity errors are fundamentally invariant with respect to the obstacle density m :

$$\frac{\|\mathbf{u} - \mathbf{u}_h\|_{L^2}}{\|\mathbf{u}\|_{L^2}} = \mathcal{O}(1), \quad \frac{|\mathbf{u} - \mathbf{u}_h|_{H^1}}{|\mathbf{u}|_{H^1}} = \mathcal{O}(1). \quad (57)$$

This theoretical derivation is corroborated by Table 2, where the relative velocity errors (Rel. $\mathbf{u}$) stabilize around $\sim 1.2\%$ for L^2 and $\sim 10\%$ for H^1 across all densities and both setups. Consequently, Table 3 shows that their convergence orders stagnate at 0.

Step 3: The Asymmetric inf-sup Pollution on Pressure. Crucially, the pressure approximation is not independent; it is contaminated by the absolute velocity error through the residual duality pairing. Bounding this yields the standard a priori estimate:

$$\|p - p_h\|_{L^2} \leq C \underbrace{\beta_h^{-1}}_{\mathcal{O}(m)} \times \left(\|\nabla(\mathbf{u} - \mathbf{u}_h)\|_{L^2} + \inf_{q_h} \|p - q_h\|_{L^2} \right). \quad (58)$$

In Setup B, enforcing a macroscopic parabolic profile triggers a severe boundary layer resonance. We observe that the absolute velocity error subsequently diverges: $\|\nabla(\mathbf{u} - \mathbf{u}_h)\|_{L^2} = \mathcal{O}(m)$ (seen in Table 2, jumping consistently from 2.76 at $m = 4$ to an extreme 83.8 at $m = 128$). Amplified by the degraded inf-sup constant $\beta_h^{-1} = \mathcal{O}(m)$, this massive microscopic error causes the absolute pressure error to grow quadratically. Because the reference pressure inherently grows as $\mathcal{O}(m^2)$ (Darcy's law), the relative pressure error completely stagnates:

$$\frac{\|p - p_h\|_{L^2}}{\|p_{\text{ref}}\|_{L^2}} \sim \frac{\mathcal{O}(m^2)}{\mathcal{O}(m^2)} = \mathcal{O}(1). \quad (59)$$

This theoretical deadlock is corroborated by Table 3, where the pressure convergence order collapses entirely to 0.03.

Surprisingly, in Setup A, the relative pressure error successfully evades this stagnation. We hypothesize that the natural Neumann condition induces a compatible macroscopic field that allows for a weak homogenization cancellation, effectively suppressing the velocity pollution term prior to the Cauchy-Schwarz bounding. However, establishing a rigorous two-scale convergence argument to fully capture this phenomenon falls beyond the primary scope of this paper, and we leave its formal mathematical resolution as an open question.

From a practical standpoint, the flow in realistic DLD devices is intrinsically driven by a fixed pressure drop. Setup B primarily serves as an idealized mathematical stress-test to demonstrate the extreme sensitivity of multiscale Stokes systems to macroscopic constraints. Therefore, we purposefully focus our discussion on the continuous stability properties and physical implications under the natural, pressure-driven regime.

Under pressure-driven (Neumann) conditions, the relative pressure error retains first-order convergence because the macroscopic field is compatible with the natural boundary formulation. Under velocity-driven (Dirichlet) conditions, the inf-sup penalty amplifies the absolute velocity error by $\mathcal{O}(m)$, matching the Darcy pressure growth and driving the relative pressure convergence rate to 0.03. The penalty is a structural incompatibility between macroscopic Dirichlet constraints and multiscale Stokes systems, not a discretization defect.

Table 2: Validation 2: Error amplification with fixed local resolution ($N_g = 6$). As obstacle density m increases, the absolute physical norms follow consistent Darcy scalings ($\|p\| \sim m^2$) in both setups. The relative velocity errors remain stationary ($\approx 1.2\% - 1.4\%$), demonstrating that the spatial discretization capacity is invariant. In contrast, the relative pressure errors suffer severe degradation under Dirichlet constraints due to the inf-sup penalty.

Setup A: Pressure-Driven (Neumann, Pressure scaled by m^2)								
m	$h_{\max}$	$\ p_{\text{ref}}\ _{L^2}$	Rel. p (L^2)	$\ \mathbf{u}_{\text{ref}}\ _{L^2}$	Abs. $\mathbf{u}$ (L^2)	Rel. $\mathbf{u}$ (L^2)	Abs. $\mathbf{u}$ (H^1)	Rel. $\mathbf{u}$ (H^1)
2	0.1074	4.97e-01	1.41e-02	5.75e-03	8.21e-05	1.43e-02	8.82e-03	1.07e-01
4	0.0549	1.91e+00	7.85e-03	7.08e-03	8.64e-05	1.22e-02	1.85e-02	9.81e-02
8	0.0275	7.60e+00	4.34e-03	7.60e-03	9.87e-05	1.30e-02	4.09e-02	1.03e-01
12	0.0184	1.72e+01	2.97e-03	7.79e-03	1.02e-04	1.31e-02	6.23e-02	1.02e-01
16	0.0137	3.08e+01	2.18e-03	7.87e-03	9.82e-05	1.25e-02	8.10e-02	9.91e-02
32	0.0069	1.24e+02	1.13e-03	8.01e-03	1.06e-04	1.32e-02	1.71e-01	1.03e-01
64	0.0034	4.98e+02	5.69e-04	8.07e-03	1.05e-04	1.31e-02	3.44e-01	1.03e-01
128	0.0017	2.00e+03	2.98e-04	8.11e-03	1.05e-04	1.30e-02	6.85e-01	1.02e-01

Setup B: Velocity-Driven Pipe Flow (Fixed Macroscopic Profile $\mathbf{u} \sim 1$)								
m	$h_{\max}$	$\ p_{\text{ref}}\ _{L^2}$	Rel. p (L^2)	$\ \mathbf{u}_{\text{ref}}\ _{L^2}$	Abs. $\mathbf{u}$ (L^2)	Rel. $\mathbf{u}$ (L^2)	Abs. $\mathbf{u}$ (H^1)	Rel. $\mathbf{u}$ (H^1)
2	0.1074	9.60e+01	1.45e-02	1.10e+00	1.53e-02	1.39e-02	1.69e+00	1.07e-01
4	0.0549	2.84e+02	8.48e-03	1.05e+00	1.25e-02	1.19e-02	2.76e+00	9.83e-02
8	0.0275	1.02e+03	5.49e-03	1.02e+00	1.29e-02	1.27e-02	5.48e+00	1.03e-01
12	0.0184	2.23e+03	4.41e-03	1.01e+00	1.29e-02	1.28e-02	8.07e+00	1.03e-01
16	0.0137	3.92e+03	3.62e-03	1.00e+00	1.22e-02	1.22e-02	1.03e+01	9.92e-02
32	0.0069	1.54e+04	3.48e-03	9.96e-01	1.29e-02	1.29e-02	2.13e+01	1.03e-01
64	0.0034	6.13e+04	3.41e-03	9.92e-01	1.27e-02	1.28e-02	4.23e+01	1.03e-01
128	0.0017	2.44e+05	3.34e-03	9.91e-01	1.26e-02	1.27e-02	8.38e+01	1.03e-01

Table 3: Empirical convergence rates for the relative errors calculated from Validation 2. As dictated by (56), the relative velocity convergence strictly stagnates at zero ($\mathcal{O}(1)$) across all resolutions up to $m = 128$. Consequently, under Setup A, the relative pressure error elegantly maintains optimal 1st-order convergence. Under Setup B, the extreme absolute velocity error, amplified by the inf-sup penalty, matches the physical Darcy pressure explosion, driving the pressure convergence rate to a complete standstill (0.03).


m	$h_{\max}$	Setup A: Pressure-Driven			Setup B: Velocity-Driven		
		Order (p)	Order ($\mathbf{u}_{L^2}$)	Order ($\mathbf{u}_{H^1}$)	Order (p)	Order ($\mathbf{u}_{L^2}$)	Order ($\mathbf{u}_{H^1}$)
2	0.1074	—	—	—	—	—	—
4	0.0549	0.87	0.23	0.12	0.80	0.23	0.12
8	0.0275	0.85	-0.08	-0.06	0.62	-0.09	-0.06
12	0.0184	0.94	-0.01	0.00	0.55	-0.02	0.00
16	0.0137	1.06	0.16	0.11	0.67	0.16	0.11
32	0.0069	0.94	-0.08	-0.05	0.05	-0.08	-0.06
64	0.0034	0.99	0.01	0.00	0.03	0.01	0.00
128	0.0017	0.94	0.01	0.01	0.03	0.01	0.01

5.4. Buffer Zones Do Not Cure the Pressure Stagnation


In practical multiscale fluid simulations, appending upstream and downstream buffer zones is a standard computational practice to mitigate boundary-induced numerical artifacts. Depending on the discipline, two distinct buffer configurations are commonly employed:

1. **Configuration 1: Patterned Buffer (Classical Oversampling).** Widely used in homogenization and multiscale finite element methods [25], this involves extending the domain using the identical periodic micro-pillar array. The premise is that high-frequency boundary layers decay exponentially within this patterned buffer.
2. **Configuration 2: Un-patterned Buffer (Realistic Empty Channel).** Standard in engineering CFD, this involves appending obstacle-free straight channels. The intuitive expectation is that the flow will naturally transition from a fully developed Poiseuille profile into the complex array.

As illustrated in Fig. 4, the computational domain is extended by a buffer zone of length $L_b = 0.5$ at both ends.



(a) Configuration 1: Patterned Buffer setup. The central domain $(0, L_x) \times (0, L_y)$ is extended by periodic buffer zones of width $L_b = 0.5$ at both ends. For $\epsilon = 1/m$, this results in a total simulation domain of $3m \times m$ unit cells.



(b) Configuration 2: Un-patterned Buffer setup. The central domain containing $2m \times m$ unit cells is augmented with obstacle-free straight channels of width $L_b = 0.5$ at the inlet and outlet to allow for flow development.

Figure 4: Schematic of the extended computational domains with different buffer configurations. The core geometry remains a DLD micro-pillar array with unit cell size $\epsilon = 1/m$.

To evaluate whether these conventional remedies can mitigate the error degradation observed under macroscopic Dirichlet constraints, we augmented Setup B by appending buffer regions of length $L_b = 0.5$ to both ends of the domain. We extract and strictly evaluate the relative errors within the central interior array (Ω_{inner}).

As documented in Table 4, while both buffering strategies successfully stabilize the velocity approximation, the relative pressure error exhibits a persistent zero-order stagnation, refusing to converge beyond $\sim 3.4 \times 10^{-3}$.

Table 4: Error amplification and empirical convergence rates under two distinct extended buffer configurations ($L_b = 0.5$). Regardless of whether the buffer is patterned (oversampling) or un-patterned (empty channel), the relative pressure error in the interior domain severely stagnates. Both convergence rates collapse toward zero as the density m increases.

m	$h_{\max}$	Config 1: Patterned Buffer (Oversampling)				Config 2: Un-patterned Buffer (Empty Channel)			
		Rel. p (L^2)	Order (p)	Rel. u (H^1)	Order (u_{H^1})	Rel. p (L^2)	Order (p)	Rel. u (H^1)	Order (u_{H^1})
2	0.107 – 0.110	1.47e-02	–	1.07e-01	–	1.58e-02	–	1.07e-01	–
4	0.055	8.56e-03	0.80	9.93e-02	0.11	9.24e-03	0.78	1.06e-01	0.01
8	0.027	5.47e-03	0.65	1.02e-01	-0.04	5.52e-03	0.74	1.04e-01	0.03
12	0.018	4.39e-03	0.54	1.03e-01	-0.02	4.45e-03	0.54	1.03e-01	0.02
16	0.014	3.63e-03	0.66	9.98e-02	0.11	3.82e-03	0.66	1.02e-01	0.04
32	0.007	3.47e-03	0.06	1.04e-01	-0.06	3.45e-03	0.15	1.03e-01	-0.01

These empirical findings serve as a necessary cautionary note for the simulation of highly heterogeneous systems. The fact that the pressure error remains locked at $\mathcal{O}(1)$ despite sufficient physical buffering presents a significant analytical challenge. The precise mathematical mechanisms driving this persistent boundary-induced stagnation, as well as the design of universally compatible boundary treatments for multiscale Stokes formulations, remain unclear within our current theoretical framework and are left as open problems for future investigation.

Neither patterned nor un-patterned buffer zones cure the pressure stagnation under macroscopic Dirichlet constraints. The error floor at $\sim 3.4 \times 10^{-3}$ is a consequence of the inf-sup penalty, not a boundary artifact that physical buffering can remove.

5.5. Algorithmic Implications: Adaptive AL Scaling

Consider the discrete saddle-point system arising from a stable mixed finite element discretization (e.g., Taylor-Hood $P_2 - P_1$ elements) of the Stokes problem on the pillar-array domain Ω_m . The conditioning of the standard discrete pressure Schur complement $S = BA^{-1}B^T$ is strictly bounded by the discrete inf-sup constant β_h , yielding $\kappa(M_p^{-1}S) \propto \beta_h^{-2}$. According to our main theorems, since the continuous constant degrades as $\beta = \Theta(m^{-1})$, any conforming discrete space inherently suffers from $\beta_h \leq \mathcal{O}(m^{-1})$, leading to $\kappa(M_p^{-1}S) \geq \mathcal{O}(m^2)$. This severe ill-conditioning causes the stagnation of standard Krylov subspace methods.

To restore the robustness of the outer iterations, the Augmented Lagrangian (AL) method introduces a grad-div stabilization term $\gamma B^T \text{diag}(M_p)^{-1} B u_h$ into the momentum equation. The modified Schur complement S_γ exhibits a significantly improved condition number bound [18]:

$$\kappa(M_p^{-1}S_\gamma) \leq C \left(1 + \frac{1}{\gamma \mu^{-1} \beta_h^2} \right), \quad (60)$$

where C is a constant independent of the mesh size h and the geometry.

In standard literature, the penalty parameter γ is often chosen heuristically. However, leveraging our pre-asymptotic bound $\beta_h \propto m^{-1}$, we can now rigorously derive the optimal strategy. To ensure that the outer Krylov solver remains completely robust and independent of the pillar density m (i.e., enforcing $\kappa(M_p^{-1}S_\gamma) = \mathcal{O}(1)$), the penalty parameter must be adaptively scaled according to:

$$\gamma = \gamma_0 \cdot m^2 \propto \epsilon^{-2}, \quad (61)$$

where $\gamma_0 > 0$ is a baseline constant of $\mathcal{O}(1)$. This strategy directly translates our continuous PDE analysis into an actionable, parameter-free stabilization strategy for microfluidic simulations.

5.6. AL Scaling Achieves Density-Independent Convergence on Square Arrays

Before extending our methodology to complex microfluidic designs, we first validate the proposed parameter-free Augmented Lagrangian (AL) preconditioner on a standard benchmark: a regular square micro-pillar array ($\delta = 0$). We maintain a constant solid packing ratio $\kappa = 0.25$ (i.e., the pillar radius is fixed at 0.25 times the characteristic unit cell diameter) and systematically refine the characteristic mesh density m from 2 to 64.

To rigorously evaluate the algorithmic scalability and preclude any false convergence artifacts (such as volumetric locking inherent to high penalty parameters in $P_2 - P_1$ elements), the FGMRES solver is driven by the unpreconditioned physical residual, enforced with a stringent absolute tolerance of 10^{-10} . Moreover, to isolate the macroscopic conditioning of the pressure Schur complement from the specific

choice of inner algebraic preconditioners, we employ a robust sparse direct solver (e.g., UMFPACK) for all inner velocity block inversions. The development of optimal, scalable multigrid preconditioners for the highly ill-conditioned AL velocity block in these densely packed geometries is a distinct and substantial challenge [18, 19], which falls outside the scope of this theoretical benchmark and is deferred to our future algorithmic research.




Figure 5: Velocity magnitude contours and streamlines of the steady Stokes flow through the square array ($\delta = 0, \kappa = 0.25$). The flow fields computed on the intermediate-density configuration (a) and the high-density configuration (b) share an identical colormap range. The results demonstrate that the highly scaled AL preconditioner accurately resolves the complex fluid tortuosity and local shear layers without exhibiting any numerical locking, even as the system approaches half a million degrees of freedom.

Table 5: Iteration count comparison for the standard square micro-pillar array ($\delta = 0, \kappa = 0.25$). The Scaled AL method effectively mitigates the ill-conditioning, achieving optimal algorithmic scalability as the degrees of freedom (DoFs) scale up to 1.85 million.

Density ($m \times m$)	Total DoFs	Std. Stokes Iters ($\gamma = 0$)	AL-Stokes Iters ($\gamma = m^2$)
2×2	1,795	34	25
4×4	7,183	53	25
8×8	29,071	86	25
16×16	116,872	154	23
32×32	468,574	277	22
64×64	1,852,843	437	22

As shown in Table 5, the standard formulation ($\gamma = 0$) suffers algorithmic degradation, with iteration counts growing from 34 to 437 as the obstacle density increases. This deterioration is a direct algebraic consequence of the degrading inf-sup stability. The Scaled AL method neutralizes this saddle-point pathology. Dictated by our strategy ($\gamma = m^2$), the AL preconditioner remains bounded and even slightly decreases, stabilizing at 22 steps for a system approaching 1.85 million DoFs. This density-independent convergence establishes a baseline for simulating more complex microfluidic topologies.

The $\gamma = m^2$ scaling reduces FGMRES iterations from 437 to 22 at 1.85M DoFs and eliminates the $\mathcal{O}(m)$ growth observed in the standard formulation. Iteration counts are bounded independent of pillar density, confirming the algebraic implication of our sharp inf-sup bound.

5.7. AL Robustness Holds Under Geometric Asymmetry ($\delta = 1/3$)

To definitively demonstrate the generalizability and practical engineering efficacy of our parameter-free stabilization strategy, we turn to a realistic Deterministic Lateral Displacement (DLD) microfluidic array. Unlike standard Cartesian grids, DLD arrays are characterized by a shifted Bravais lattice ($\delta > 0$), which breaks geometric symmetry and introduces complex, tortuous streamlines responsible for particle fractionation.

We construct a sequence of DLD pillar arrays with a characteristic row shift fraction $\delta = 1/3$ and maintain the solid geometric ratio $\kappa = 0.25$. The mesh density m is scaled from 2 up to an extreme case of 64, culminating in a highly ill-conditioned saddle-point system with over 1.76 million degrees of freedom. Crucially, to preclude any spurious convergence artifacts (e.g., false volumetric locking) that can be induced by the highly scaled penalty parameter ($\gamma = m^2$) in the $P_2 - P_1$ space, we strictly evaluate the unpreconditioned physical residual. The FGMRES solver is enforced to satisfy a stringent absolute tolerance of 10^{-10} alongside a relative tolerance of 10^{-8} .




Figure 6: Velocity magnitude contours and streamlines of the steady Stokes flow through the DLD array ($\delta = 1/3$, $\kappa = 0.25$). The flow fields computed on the intermediate-density configuration (a) and the high-density configuration (b) share an identical colormap range. The results demonstrate that the highly scaled AL preconditioner accurately resolves the complex fluid tortuosity and local shear layers without exhibiting any numerical locking, even as the system approaches half a million degrees of freedom.

Table 6: Comparison of FGMRES outer iterations for a DLD array with a row shift fraction $\delta = 1/3$. Even with broken symmetry, highly tortuous flow paths, and stringent convergence criteria (10^{-10} absolute tolerance), the AL method exhibits robust algorithmic scalability.

Density ($m \times m$)	Total DoFs	Std. Stokes Iters ($\gamma = 0$)	AL-Stokes Iters ($\gamma = m^2$)
2×2	1,755	33	24
4×4	6,860	57	28
8×8	27,467	95	27
16×16	109,983	175	27
32×32	441,508	324	26
64×64	1,766,167	687	24

Table 6 shows that the geometric asymmetry and tortuous flow paths of the DLD device exacerbate the algebraic failure of the standard decoupled formulation. Compared to the square array, the standard Stokes iterations degrade more aggressively, from 33 to 687 iterations as pillar density increases. The Augmented Lagrangian method, governed by $\gamma = m^2$, remains bounded under this added geometric complexity, stabilizing at 24 iterations for the 1.76M-DoF system. This experiment corroborates our framework: a parameter-free stabilization strategy insensitive to both the obstacle density and the geometric configurations of modern microfluidic simulations.

The AL scaling is insensitive to geometric asymmetry. Iteration counts remain bounded at 24 on a 1.77M-DoF DLD array, while the standard formulation degrades to 687. The $\gamma \propto m^2$ rule addresses the geometric ill-conditioning itself, not a symmetry-specific property of the square lattice.

6. Conclusions and Future Directions

In this paper, we have rigorously established a sharp pre-asymptotic estimate for the continuous inf-sup constant in densely packed microfluidic arrays. We proved that the constant deteriorates at a

sharp rate of $\Theta(m^{-1})$ as the pillar density m increases, exposing the fundamental geometric penalty that dictates the stability of the multiscale Stokes system.

Our numerical investigations map this stability degradation to its two concrete manifestations. On the discretization side, the sharp bound enters Brezzi's a priori estimate through the factor $\beta_h^{-1} = \mathcal{O}(m)$, and we observe an asymmetric pollution under macroscopic boundary conditions: Neumann pressure-drop inflows retain first-order pressure convergence, while Dirichlet inflow profiles drive the relative pressure error to a complete $\mathcal{O}(1)$ stagnation that buffer zones do not cure. On the algebraic side, the Schur complement condition scales as $\mathcal{O}(m^2)$, dictating the stagnation of standard iterative methods. Leveraging the sharp bound, we proposed an adaptively scaled Augmented Lagrangian (AL) method with $\gamma \propto m^2$. Experiments on square and asymmetric DLD arrays show that this parameter-free scaling neutralizes the geometric ill-conditioning and achieves density-independent convergence for systems with millions of degrees of freedom.

Future research will focus on developing geometry-aware multigrid preconditioners for the inner AL velocity blocks and extending this framework to multi-scale Fluid-Structure Interaction (FSI) problems. These advancements will enable predictive, fully resolved simulations of complex cell deformations and trajectories within intricate micro-channel topologies.

References

- [1] L. R. Huang, E. C. Cox, R. H. Austin, J. C. Sturm, Continuous particle separation through deterministic lateral displacement, *Science* 304 (5673) (2004) 987–990. doi:10.1126/science.1094567.
- [2] J. McGrath, M. Jimenez, H. Bridle, Deterministic lateral displacement for particle separation: a review, *Lab Chip* 14 (21) (2014) 4139–4158. doi:10.1039/c4lc00939h.
URL <http://dx.doi.org/10.1039/C4LC00939H>
- [3] D. W. Inglis, J. A. Davis, R. H. Austin, J. C. Sturm, Critical particle size for fractionation by deterministic lateral displacement, *Lab on a Chip* 6 (5) (2006) 655. doi:10.1039/b515371a.
URL <http://dx.doi.org/10.1039/B515371A>
- [4] B. M. Dincau, A. Aghilinejad, T. Hammersley, X. Chen, J.-H. Kim, Deterministic lateral displacement (DLD) in the high reynolds number regime: high-throughput and dynamic separation characteristics, *Microfluidics and Nanofluidics* 22 (6) (May 2018). doi:10.1007/s10404-018-2078-9.
URL <http://dx.doi.org/10.1007/s10404-018-2078-9>
- [5] A. Meier, E. Bänsch, F. Frank, Schur preconditioning of the Stokes equations in channel-dominated domains, *Computer Methods in Applied Mechanics and Engineering* 398 (2022) 115264. doi:10.1016/j.cma.2022.115264.
URL <http://dx.doi.org/10.1016/j.cma.2022.115264>
- [6] D. Boffi, F. Brezzi, M. Fortin, *Mixed finite element methods and applications*, Vol. 44, Springer, 2013. doi:10.1007/978-3-642-36519-5.
- [7] K.-A. Mardal, R. Winther, Preconditioning discretizations of systems of partial differential equations, *Numerical Linear Algebra with Applications* 18 (1) (2010) 1–40. doi:10.1002/nla.716.
URL <http://dx.doi.org/10.1002/nla.716>
- [8] D. Loghin, A. J. Wathen, Analysis of preconditioners for saddle-point problems, *SIAM Journal on Scientific Computing* 25 (6) (2004) 2029–2049. doi:10.1137/S1064827502418203.
- [9] E. V. Chizhonkov, M. A. Olshanskii, On the domain geometry dependence of the LBB condition, *ESAIM: Mathematical Modelling and Numerical Analysis* 34 (5) (2000) 935–951. doi:10.1051/m2an:2000110.
URL <http://dx.doi.org/10.1051/m2an:2000110>
- [10] M. Dobrowolski, On the LBB constant on stretched domains, *Mathematische Nachrichten* 254-255 (1) (2003) 64–67. doi:10.1002/mana.200310052.
URL <http://dx.doi.org/10.1002/mana.200310052>
- [11] E. Sande, T. Koch, M. Kuchta, K.-A. Mardal, On a robust inf-sup condition for the Stokes problem in slender domains—with application to preconditioning, *arXiv preprint arXiv:2510.24590* (2025).

- [12] F. Brezzi, On the existence, uniqueness and approximation of saddle-point problems arising from Lagrangian multipliers, *Revue française d'automatique, informatique, recherche opérationnelle. Analyse numérique* 8 (R2) (1974) 129–151. doi:10.1051/m2an/197408r201291.
URL <http://dx.doi.org/10.1051/m2an/197408R201291>
- [13] V. Girault, P.-A. Raviart, *Finite element methods for Navier-Stokes equations: theory and algorithms*, Springer Science & Business Media, 2012.
- [14] A. Ern, J.-L. Guermond, *Theory and practice of finite elements*, Vol. 159, Springer, 2004.
- [15] C. Bacuta, A unified approach for Uzawa algorithms, *SIAM Journal on Numerical Analysis* 44 (6) (2006) 2633–2649. doi:10.1137/050630714.
URL <http://dx.doi.org/10.1137/050630714>
- [16] G. Allaire, Homogenization of the Navier–Stokes equations in open sets perforated with tiny holes ii: Non-critical sizes of the holes for a volume distribution and a surface distribution of holes, *Archive for Rational Mechanics and Analysis* 113 (3) (1991) 261–298. doi:10.1007/bf00375066.
URL <http://dx.doi.org/10.1007/BF00375066>
- [17] Y. Lu, Homogenization of Stokes equations in perforated domains: A unified approach, *Journal of Mathematical Fluid Mechanics* 22 (3) (Aug. 2020). doi:10.1007/s00021-020-00506-6.
URL <http://dx.doi.org/10.1007/s00021-020-00506-6>
- [18] M. Benzi, M. A. Olshanskii, An augmented Lagrangian-based approach to the Oseen problem, *SIAM Journal on Scientific Computing* 28 (6) (2006) 2095–2113. doi:10.1137/050646421.
URL <http://dx.doi.org/10.1137/050646421>
- [19] P. E. Farrell, L. Mitchell, F. Wechsung, An augmented Lagrangian preconditioner for the 3d stationary incompressible Navier–Stokes equations at high Reynolds number, *SIAM Journal on Scientific Computing* 41 (5) (2019) A3073–A3096. doi:10.1137/18m1219370.
URL <http://dx.doi.org/10.1137/18M1219370>
- [20] M. Fortin, R. Glowinski, *Augmented Lagrangian methods: applications to the numerical solution of boundary-value problems*, Vol. 15, Elsevier, 2000.
- [21] R. Glowinski, P. Le Tallec, *Augmented Lagrangian and operator-splitting methods in nonlinear mechanics*, SIAM, 1989.
- [22] L. Chen, A simple construction of a Fortin operator for the two dimensional Taylor–Hood element, *Computers and Mathematics with Applications* 68 (10) (2014) 1368–1373. doi:10.1016/j.camwa.2014.09.003.
URL <http://dx.doi.org/10.1016/j.camwa.2014.09.003>
- [23] S. C. Brenner, L. R. Scott, *The Mathematical Theory of Finite Element Methods*, 3rd Edition, Vol. 15 of Texts in Applied Mathematics, Springer, 2008. doi:10.1007/978-0-387-75934-0.
- [24] R. E. Bank, A. Weiser, Some a posteriori error estimators for elliptic partial differential equations, *Mathematics of Computation* 44 (170) (1985) 283–301. doi:10.1090/s0025-5718-1985-0777265-x.
URL <http://dx.doi.org/10.1090/s0025-5718-1985-0777265-x>
- [25] T. Y. Hou, X.-H. Wu, A multiscale finite element method for elliptic problems in composite materials and porous media, *Journal of Computational Physics* 134 (1) (1997) 169–189. doi:10.1006/jcph.1997.5682.
URL <http://dx.doi.org/10.1006/jcph.1997.5682>