Peer reviewed Journals
seron press

RESEARCH

Thermodynamic-Metric Preconditioning for Entropy-Stable Implicit Simulation of Gas-Liquid Two-Fluid Networks

Lucas Tavares[1] and Bruno Azevedo[2]

[1]Universidade Federal do Vale do São Francisco, Avenida José de Sá Maniçoba S/N, Centro, Petrolina, Brazil
[2]Universidade Estadual de Santa Cruz, Rodovia Jorge Amado Km 16, Salobrinho, Ilhéus, Brazil

Abstract

Gas-liquid multiphase transport in annular channels and networked conduits is often modeled by two-fluid balance laws whose numerical solution becomes challenging under stiff interfacial exchange, strong slip, and sparse operational constraints. In practical regimes, reliable computation requires more than accuracy: it requires positivity of phase masses and volume fractions, bounded thermodynamic states, and dissipation consistent with the second law. This paper develops a structure-preserving implicit simulation framework for one-dimensional two-fluid networks based on a thermodynamic metric that acts simultaneously as an entropy symmetrizer and as a preconditioner for Newton–Krylov solves. The central contribution is a preconditioning strategy defined in entropy-variable coordinates, where the Jacobian of the discretized residual is rendered close to symmetric in a weighted inner product and where stiff source couplings are balanced by metric-consistent scaling. A complementary constrained projection is introduced to enforce admissibility during nonlinear iterations without destroying differentiability, enabling stable sensitivity analysis through discrete adjoints. The method integrates network junction constraints through a Schur-complement coupling that preserves conservation while maintaining the entropy inequality at the fully discrete level. Numerical experiments on annular segments and branched networks demonstrate that the thermodynamic-metric preconditioner reduces Krylov iterations across wide operating envelopes and suppresses solver stagnation near phase-disappearance limits, while the admissibility projection eliminates nonphysical intermediate iterates that would otherwise destabilize implicit solves. The resulting approach targets robust, scalable simulation and gradient computation for multiphase networks under stiff closures and limited numerical damping.

1. Introduction

Implicit simulation of gas–liquid two-phase flows remains a persistent technical challenge because the governing models combine hyperbolic transport with stiff interfacial and wall exchange processes, strong nonlinear thermodynamics, and constraint sets that reflect physical admissibility rather than mathematical convenience [1]. In annular and near-annular conditions, the state can approach phase-disappearance limits locally, slip can change rapidly, and closures for momentum and energy exchange can introduce effective time scales far smaller than those implied by bulk convection. Explicit time integration then becomes impractical, while fully implicit methods can fail to converge or can converge to nonphysical states unless the nonlinear solve is carefully stabilized. The difficulty is amplified in network settings, where junction constraints couple segments algebraically and where boundary actuation can abruptly shift operating points. These factors motivate solver designs that are not merely fast on average but robust with respect to extreme local states, closure stiffness, and the geometry of the admissible set.

A common practical response is to reduce model fidelity, add numerical diffusion, or clip state variables when they drift out of bounds. Such responses can be operationally effective in narrow contexts, yet they obscure the underlying stability structure and complicate sensitivity analysis and optimization, where differentiability and consistency matter. If an implicit method relies on ad hoc clipping, the discrete residual no longer corresponds to a well-defined variational or balance principle, and gradients become unreliable. Conversely, if one insists on strict conservation and thermodynamic consistency but uses generic Newton–Krylov machinery without structure-aware scaling, the linearized systems can become severely ill conditioned, leading to stagnation or divergence precisely in the regimes where implicit simulation is most needed. The central question addressed in this paper is therefore how to design an implicit solver for two-fluid networks that is simultaneously structure preserving, robust near admissibility boundaries, and computationally scalable.

The approach adopted here is to treat the two-fluid model not as an arbitrary nonlinear system but as a dis-

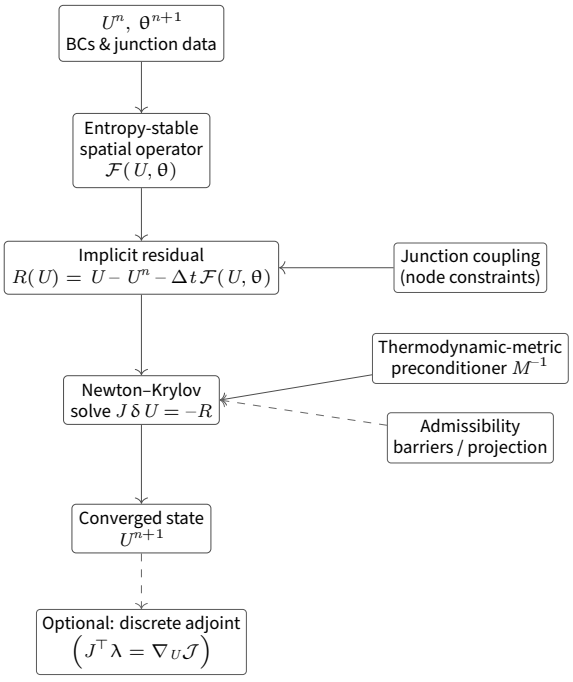


Figure 1. End-to-end architecture for entropy-stable implicit simulation of a gas–liquid two-fluid network. Conservative states advance by solving a fully implicit residual with Newton–Krylov iterations, accelerated by a thermodynamic-metric preconditioner and safeguarded by smooth admissibility enforcement that keeps phase masses, volume fractions, and thermodynamic inversions well-defined during nonlinear iterations. Junction constraints enter as coupled residual blocks so conservation is enforced at nodes while edge dynamics are advanced consistently.

Table 1. Conservative state components for the one-dimensional two-fluid model.

Component	Symbol	Definition
Gas mass per unit length	m_g	$m_g = \alpha_g \rho_g A$
Liquid mass per unit length	m_ℓ	$m_\ell = \alpha_\ell \rho_\ell A$
Gas momentum per unit length	p_g	$p_g = m_g u_g$
Liquid momentum per unit length	p_ℓ	$p_\ell = m_\ell u_\ell$
Gas total energy per unit length	e_g	$e_g = m_g E_g$
Liquid total energy per unit length	e_ℓ	$e_\ell = m_\ell E_\ell$

sipative balance law that admits an entropy function, with associated entropy variables that symmetrize the hyperbolic part and encode a natural thermodynamic metric [?]. In classical theory, entropy variables provide a coordinate system in which flux Jacobians become symmetric and where energy stability can be analyzed through convexity. In practice, however, most industrial two-fluid solvers operate in conservative variables because they are directly tied to conservation constraints, and because entropy-variable formulations can be awkward when closures are expressed in primitive variables. This paper bridges that gap by using entropy variables as an internal computational coordinate for preconditioning and for projection, while retaining a conservative-variable residual as the primary object solved to ensure conservation and consistent junction coupling.

A second motivating observation is that, while regime labels are not themselves a numerical method, the success of data-driven regime recognition in annular channels suggests that multiphase state trajectories occupy structured subsets of state space and that measurable signals can strongly constrain which parts of that subset are realized. Manikonda et al. (2021) described early machine-learning classification models that, in annular flow-channel experiments, could infer gas–liquid flow regimes with very high predictive performance and operationally low latency, indicating that annular multiphase behavior produces consistent signatures in accessible measurements rather than arbitrary variability [2]. The present paper does not build a classifier and does not frame the solver around regime labeling, but this empirical observation reinforces a methodological stance: numerical methods should be designed to respect the geometry of physically realized states, including admissibility boundaries and dissipative struc-

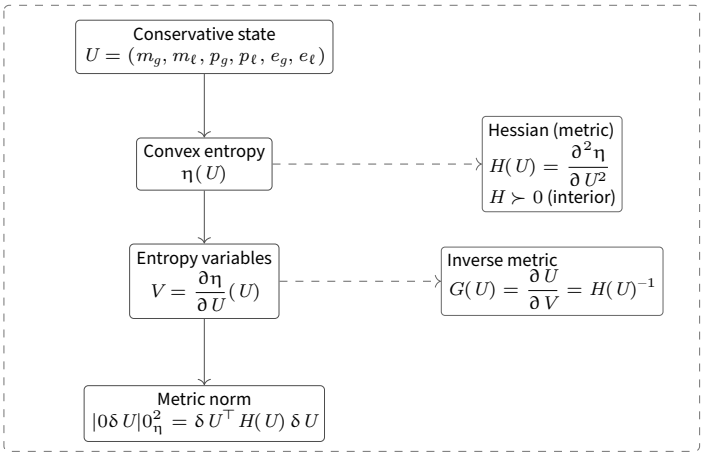


Figure 2. Entropy geometry used as a computational coordinate system. A convex entropy defines entropy variables and a positive-definite Hessian that induces a thermodynamic metric. The inverse metric provides a structured scaling between conservative perturbations and entropy-variable perturbations, enabling metric-aware linear solves and constraint handling in a thermodynamically meaningful norm.

Table 2. Primitive and thermodynamic quantities derived from the conservative variables.

Quantity	Symbol	Relation
Gas volume fraction	α_g	$\alpha_g = m_g/(\rho_g A)$
Liquid volume fraction	α_ℓ	$\alpha_\ell = m_\ell/(\rho_\ell A)$
Phase velocity	u_k	$u_k = p_k/m_k$
Specific total energy	E_k	$E_k = e_k^{\mathrm{int}} + \frac{1}{2}u_k^2$
Temperature	T_k	From EOS: $T_k = T(\rho_k, e_k^{\mathrm{int}})$
Pressure	p	From EOS: $p = p(\rho_k, e_k^{\mathrm{int}})$

ture, rather than to rely on generic algebraic robustness mechanisms that ignore that geometry.

A further motivation is the increasing demand for differentiable simulation. Whether the goal is parameter calibration, uncertainty quantification, design optimization, or control synthesis, gradients through multiphase solvers are becoming as important as forward predictions. Differentiable simulation is particularly fragile in stiff nonlinear systems because small perturbations can lead to large changes in iteration counts, active constraints, and local linearization quality. A solver architecture that is stable and that enforces admissibility through smooth mechanisms therefore directly supports reliable sensitivity analysis. This paper explicitly targets this requirement by constructing preconditioning and projection operations that are compatible with discrete adjoints, enabling stable gradient computation even when the forward solve operates near phase limits [3].

The central technical contribution is a thermodynamic-metric preconditioner for Newton–Krylov implicit solves of discretized two-fluid network equations. The preconditioner is derived from an entropy-variable mapping and can be interpreted as a Riemannian metric on the state manifold induced by convex entropy. In that metric, the linearized residual becomes closer to symmetric and better balanced across conserved components, while stiff source couplings become naturally scaled by thermodynamic sensitivities. This reduces the spread in Jacobian eigenvalues and improves Krylov convergence. The preconditioner is complemented by an admissibility projection that enforces positivity and a discrete entropy inequality during nonlinear iterations. Unlike hard clipping, the projection is posed as a constrained minimization in entropy-variable space, yielding a smooth correction that can be differentiated through. Junction coupling is handled by a Schur-complement construction that respects conservation at nodes and that preserves the entropy inequality when edges are coupled.

The paper is organized around these elements. The next section formalizes a two-fluid network discretization, defines an entropy function and associated entropy variables, and introduces the thermodynamic metric that will be used for preconditioning and projection. The following section derives the implicit time integration scheme and the Newton–Krylov solve, then constructs the thermodynamic-metric preconditioner and analyzes its effect on linear system conditioning [4]. The next section develops the admissibility projection and its integration into nonlinear iterations, emphasizing differentiability and the preservation of a dis-

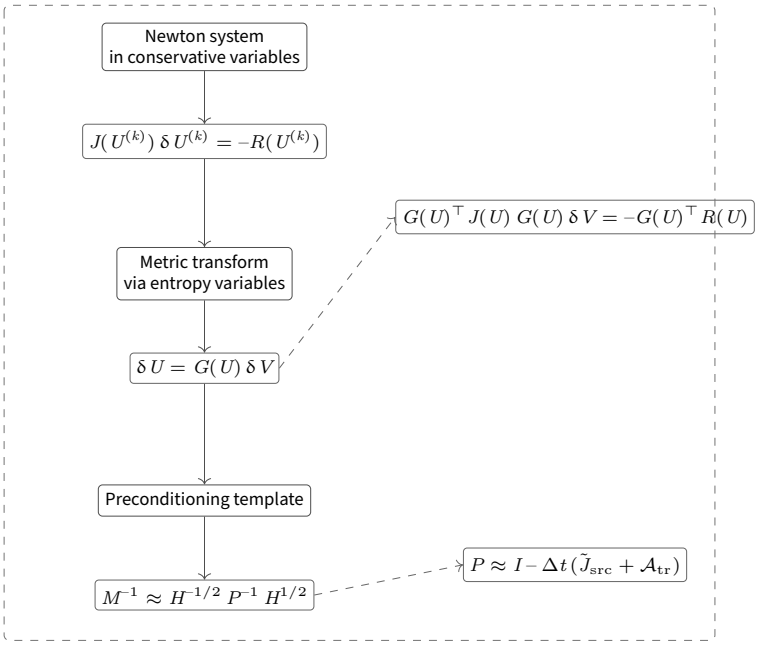


Figure 3. Metric-scaled Newton linearization. The update is expressed through the inverse entropy Hessian, converting conservative increments into entropy-variable increments. In the metric-induced inner product, the transformed operator becomes better balanced and typically closer to symmetric, motivating preconditioners that combine entropy-based scaling with sparse approximations of transport and stiff source couplings.

Table 3. Entropy-related objects and their roles in the thermodynamic metric.

Quantity	Symbol	Role
Entropy density	$\eta(U)$	Convex potential defining admissible thermodynamic states
Entropy variables	$V(U)$	$V = \partial\eta/\partial U$ (coordinate transform)
Entropy Hessian	$H(U)$	$H = \partial^2\eta/\partial U^2$, positive definite in interior
Thermodynamic metric	$\langle \delta U, \delta U\rangle_\eta$	$\delta U^\top H(U)\,\delta U$
Entropy flux	$\psi(U)$	Satisfies continuous entropy balance with dissipation

crete entropy inequality. The subsequent section extends the method to network junction coupling, presenting a Schur-complement strategy that integrates node constraints without destroying edge-level entropy stability. The final technical section reports computational experiments that evaluate solver robustness, iteration counts, and gradient stability on annular segments and networked configurations across stiff closure regimes. The conclusion summarizes implications and limitations, focusing on how thermodynamic geometry can be used as a computational tool for robust multiphase implicit simulation.

2. Entropy Variables, Thermodynamic Metrics, and Discretized Two-Fluid Residuals

A robust implicit method for two-fluid models must begin from a discretization that makes explicit both conservation structure and the dissipative constraints that define admissibility. Consider a one-dimensional network edge parameterized by an axial coordinate x with a cross-sectional area $A(x)$ and hydraulic geometry appropriate to an annulus. Let $k \in \{g, \ell\}$ denote gas and liquid. In a conservative representation, one evolves phase masses, phase momenta, and phase total energies per unit axial length. A representative conservative state on a cell is

$$U = \big(m_g,\ m_\ell,\ p_g,\ p_\ell,\ e_g,\ e_\ell\big), \tag{1}$$

where $m_k = \alpha_k \rho_k A$ are phase masses per unit length, $p_k = m_k u_k$ are phase momenta per unit length, and $e_k = m_k E_k$ are phase total energies per unit length. The variables α_k are volume fractions, ρ_k are densities, u_k are velocities, and $E_k = e_k^{\mathrm{int}} + \frac{1}{2}u_k^2$ are specific total energies, with e_k^{int} internal energy. The mapping between (m_k, e_k, p_k) and primitive variables involves an equation of state and geometric relations between α_k and film thickness-like variables in annular conditions. The details of the equation of state are not the focus of this paper, but it is assumed that a convex thermodynamic entropy exists and that pressure and temperature can be computed consistently from internal energy and density [?].

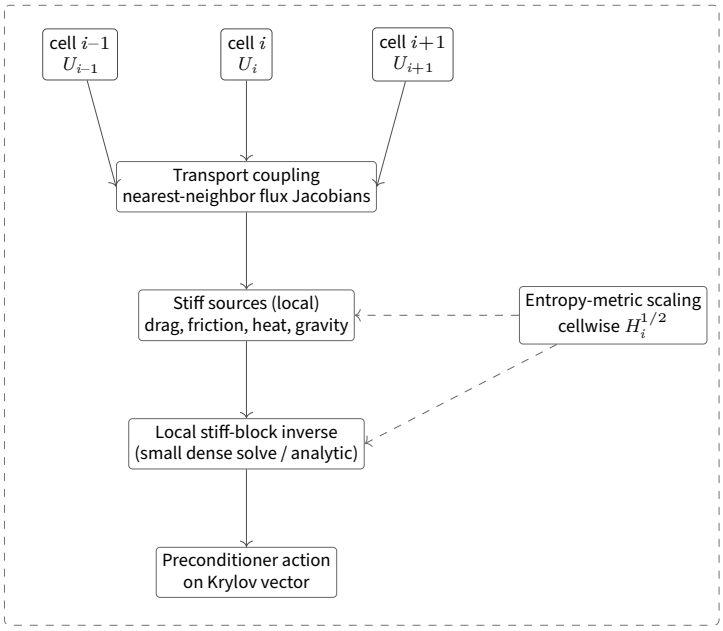


Figure 4. Cell-local structure exploited by the thermodynamic-metric preconditioner. Nearest-neighbor transport contributions enter through flux Jacobians, while stiff interfacial and wall exchange terms are treated as compact local blocks whose inversion can be approximated cheaply. Entropy-metric scaling is applied cellwise so that momentum and energy couplings remain well balanced even when a phase mass becomes small.

Table 4. Residual and Jacobian structure for fully implicit time integration.

Object	Symbol	Expression
Time-step residual	$R(U^{n+1})$	$U^{n+1} - U^n - \Delta t\,\mathcal{F}(U^{n+1}, \theta^{n+1})$
Transport Jacobian	J_{tr}	$\partial(\text{flux differences})/\partial U$
Source Jacobian	J_{src}	$\partial S(U,\theta)/\partial U$
Full Jacobian	J	$J = I - \Delta t\,(J_{\mathrm{tr}} + J_{\mathrm{src}})$
Newton correction	δU	Solves $J(U^{(k)})\,\delta U^{(k)} = -R(U^{(k)})$

Spatial discretization partitions each edge into cells indexed by i with length Δx_i. A semidiscrete conservative update for the hyperbolic transport can be written as a flux difference plus source contributions. Let $F_{i+1/2}(U)$ denote a numerical flux at the interface between cells i and $i+1$. Let $S_i(U,\theta)$ denote source terms, including gravity, wall friction, interfacial momentum exchange, and heat transfer, parameterized by closure coefficients collected in θ. The semidiscrete evolution is

$$\frac{d}{dt}U_i = -\frac{1}{\Delta x_i}\Big(F_{i+1/2} - F_{i-1/2}\Big) + S_i(U,\theta) + J_i(U,\theta), \tag{2}$$

where J_i captures junction coupling for boundary-adjacent cells. A fully implicit time integration over a time step Δt produces a nonlinear residual equation for U^{n+1} of the general form

$$R(U^{n+1}) = U^{n+1} - U^n - \Delta t\,\mathcal{F}(U^{n+1}, \theta^{n+1}) = 0, \tag{3}$$

where $\mathcal{F}$ denotes the spatial operator including flux differences, sources, and junction coupling. In a network, U denotes the concatenated vector over all edges and selected node variables.

The admissible set for U includes positivity constraints $m_k \geq 0$, bounds on volume fractions $0 \leq \alpha_k \leq 1$, and thermodynamic bounds ensuring positive temperatures and physically meaningful pressures. In two-fluid systems, these constraints are not independent; for example, α_k depends on m_k and ρ_k, and ρ_k depends on equation-of-state inversion from internal energy. It is computationally convenient to treat admissibility through a convex entropy. Let $\eta(U)$ denote an entropy density that is convex in the conservative variables over the admissible set [5]. The exact expression depends on the thermodynamic model, but the crucial property is convexity. An associated entropy flux $\psi(U)$ exists for the hyperbolic part, and dissipative sources satisfy an entropy production inequality. At the continuous level, one expects

$$\partial_t \eta(U) + \partial_x \psi(U) \leq \sigma(U,\theta), \tag{4}$$

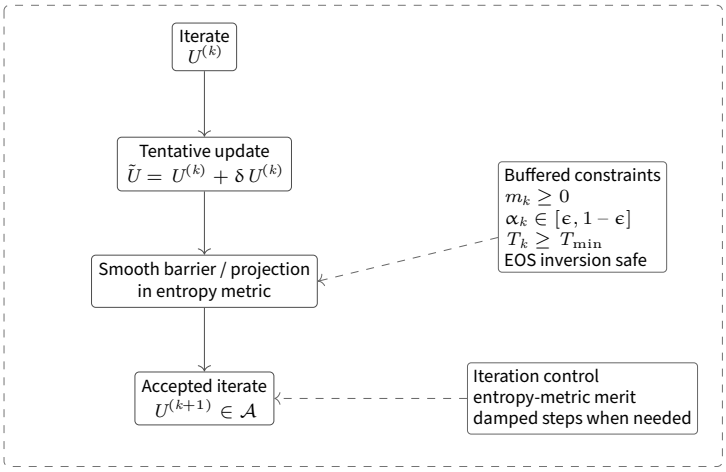


Figure 5. Admissibility enforcement embedded in nonlinear iterations. A tentative Newton update is passed through a smooth barrier/projection mechanism formulated in the thermodynamic metric, maintaining positivity and bounded phase fractions while keeping equation-of-state inversions and closure evaluations well conditioned. The same metric supports step control that discourages iterates from approaching phase-disappearance limits too aggressively.

Table 5. Thermodynamic-metric preconditioning operators in entropy-variable coordinates.

Operator	Symbol	Description
Entropy Hessian	$H(U)$	Local metric tensor on the conservative state manifold
Inverse metric	$G(U)$	$G = \partial U/\partial V = H^{-1}$
Metric-scaled Jacobian	$\tilde{J}(U)$	$\tilde{J} = G^{\top} J G$
Source contribution	$\tilde{J}_{\mathrm{src}}$	$\tilde{J}_{\mathrm{src}} = H^{1/2} J_{\mathrm{src}} H^{1/2}$
Preconditioner	M^{-1}	$M^{-1} = H^{-1/2} P^{-1} H^{1/2}$ with $P \approx H^{1/2} J H^{1/2}$

with σ nonpositive in a sign convention where entropy is convex and dissipation decreases η. At the discrete level, entropy stability is encoded by choosing numerical fluxes and source discretizations such that a discrete inequality holds. The present paper uses this structure as a computational asset rather than as a purely analytical property.

Entropy variables are defined as the gradient of entropy with respect to conservative variables:

$$V = \frac{\partial \eta}{\partial U}(U). \tag{5}$$

Convexity implies that the Hessian $H(U) = \partial^2 \eta / \partial U^2$ is positive definite in the interior of the admissible set, and that the mapping between U and V is invertible. This Hessian defines a natural inner product and thus a thermodynamic metric:

$$\langle \delta U, \delta U \rangle_{\eta} = \delta U^{\top} H(U)\, \delta U. \tag{6}$$

In the context of hyperbolic conservation laws, entropy variables symmetrize the flux Jacobians: if $A(U) = \partial f / \partial U$ is the flux Jacobian for a physical flux f, then $H(U)A(U)$ is symmetric under suitable assumptions. While two-fluid models with complex closures may not be globally symmetrizable in the same strict sense as single-fluid Euler equations, the entropy-variable mapping still provides a meaningful scaling for the conservative components, reflecting thermodynamic sensitivities and balancing the relative magnitudes of mass, momentum, and energy perturbations [6]. This scaling becomes particularly relevant for implicit solves, where ill conditioning often arises from disparate units and stiffness across components.

The thermodynamic metric is used in this paper to build a preconditioner for the Jacobian of the implicit residual. To do that, one must define the Jacobian of the residual in conservative variables and express how it transforms under the entropy-variable mapping. Let $R(U)$ denote the global residual for a time step. The Jacobian is

$$J(U) = \frac{\partial R}{\partial U}(U). \tag{7}$$

A standard Newton update solves

$$J(U^{(k)})\, \delta U^{(k)} = -R(U^{(k)}), \tag{8}$$

then updates $U^{(k+1)} = U^{(k)} + \delta U^{(k)}$. In practice, J is large and sparse and is solved approximately by Krylov methods. Krylov convergence depends strongly on the conditioning and spectrum of J. Because J includes contributions from both transport and stiff sources, its spectrum can be wide. Moreover, the scaling of conser-

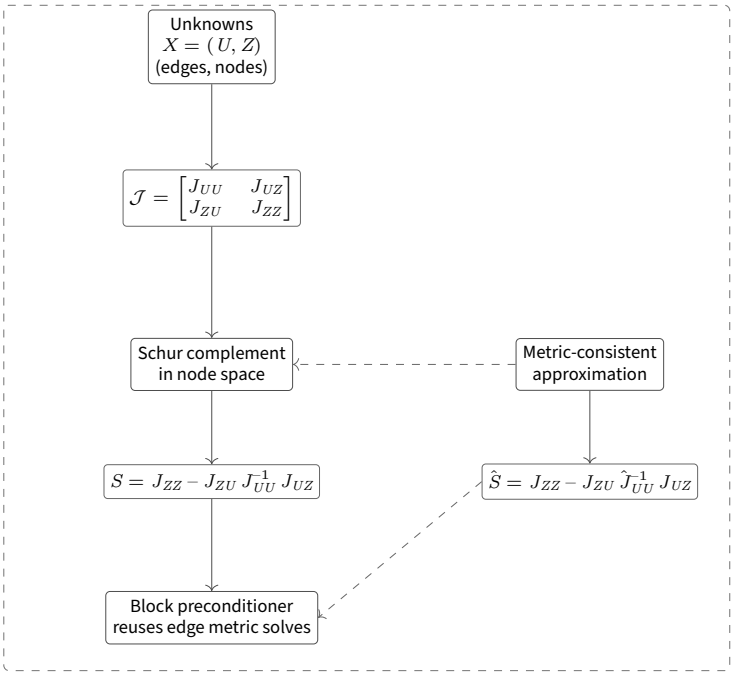


Figure 6. Network junction coupling integrated through a Schur-complement strategy. Edge states and node variables form a block Jacobian; eliminating edge unknowns yields a compact node-level operator capturing global coupling. Using the thermodynamic-metric edge preconditioner as an approximate inverse produces a metric-consistent Schur complement that improves convergence in connected networks while preserving conservation in the converged solution.

Table 6. Representative admissibility constraints for the conservative state and entropy.

Variable	Constraint	Type	Notes
Phase mass m_k	$m_k \geq 0$	Positivity	Prevents loss of phase and EOS breakdown
Volume fraction α_k	$\alpha_k \in [\epsilon, 1-\epsilon]$	Box constraint	Buffers phase-disappearance limits
Temperature T_k	$T_k \geq T_{\min}$	Lower bound	Avoids unphysical low-temperature states
Internal energy e_k^{int}	$e_k^{\text{int}} \in [e_{\min}, e_{\max}]$	Box constraint	Enforces thermodynamic consistency
Discrete entropy	$\eta(U^{n+1}) \leq \eta(U^n) + \Delta t\, \Sigma$	Inequality	Encodes step-level entropy dissipation

vative variables leads to poorly balanced block structure, particularly when phase masses are small and momenta and energies remain finite, as occurs near phase disappearance [7]. These factors motivate transforming the linear system into entropy-variable coordinates.

Under a change of variables $V = V(U)$, a Newton step can be written in terms of δV:

$$\delta U = G(U)\, \delta V, \tag{9}$$

where $G(U) = \partial U / \partial V$ is the inverse Hessian, $G(U) = H(U)^{-1}$ when $V = \partial \eta / \partial U$. Substituting into the Newton system yields

$$J(U)\, G(U)\, \delta V = -R(U). \tag{10}$$

If one left-multiplies by $G(U)^\top$, one obtains a symmetrized form in the metric induced by G:

$$G(U)^\top J(U)\, G(U)\, \delta V = -G(U)^\top R(U). \tag{11}$$

The matrix $G^\top JG$ is not exactly symmetric in general, but it is typically better balanced and closer to symmetric in regimes where entropy variables align with the dominant stiffness mechanisms. This motivates using G as a preconditioning operator. Instead of solving in δV exactly, one uses the metric structure to define a preconditioner M that approximates J^{-1} in a way that respects thermodynamic scaling.

A second role of the entropy metric is to define a smooth projection onto the admissible set. When Newton iterates stray near or outside admissibility, the equation-of-state inversion can fail or can produce nonphysical pressures. A projection defined in the entropy metric penalizes deviations in thermodynamically meaningful units rather than in raw conservative units. This can prevent the projection from introducing artificial energy changes that destabilize entropy stability. The projection will be developed later, but its conceptual foundation is the same: entropy provides a convex geometry that can be used both to guide linear solves and to enforce constraints [8].

In two-fluid models, closures introduce additional

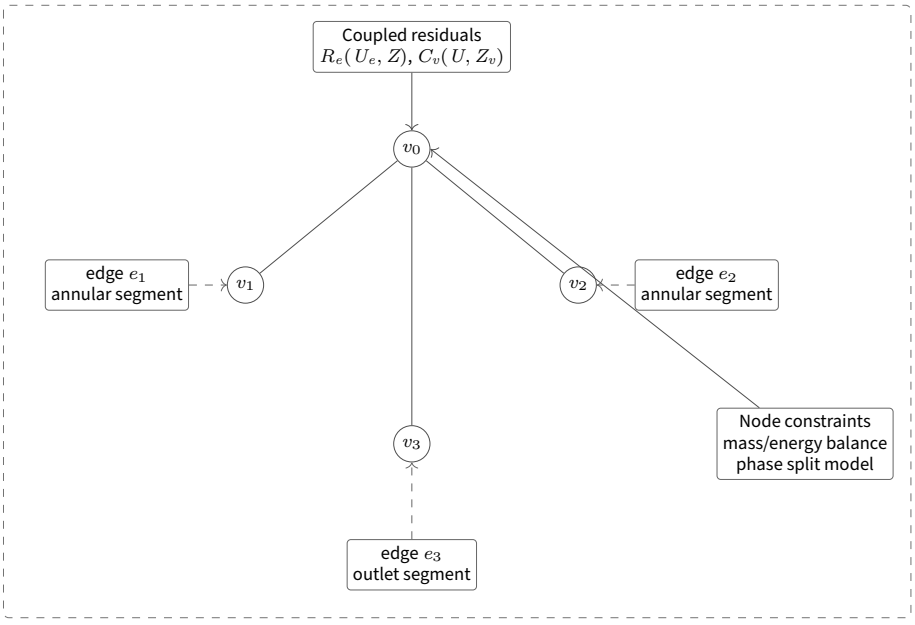


Figure 7. Representative branched two-fluid network used to motivate junction-aware solver design. Edges carry conservative two-fluid states while nodes impose algebraic constraints enforcing conservation and modeling phase splitting and pressure losses. The implicit solve couples edge and node unknowns so that global operating changes at junctions are handled without breaking entropy stability or conservation.

Table 7. Entropy-metric projection ingredients for admissibility-preserving Newton iterations.

Feature	Symbol / expression	Purpose
Projection operator	$\Pi(\tilde{U})$	Maps tentative state to admissible $U = \Pi(\tilde{U})$
Reference state	$U^\star$	Anchor for local metric in the quadratic objective
Metric-weighted distance	$(U-\tilde{U})^\top H(U^\star)(U-\tilde{U})$	Objective minimized under constraints
Barrier parameter	μ	Scales smooth penalties for constraint proximity
Merit function	$\phi(U)$	Combines residual norm with entropy-based barrier terms

complexity because sources depend on primitive variables and on regime-dependent coefficients. The implicit residual includes source terms that may be stiff, such as interfacial drag modeled as a relaxation toward equal velocities or slip relations. Such sources can be written abstractly as

$$S(U,\theta) = \mathcal{S}\big(W(U),\theta\big), \tag{12}$$

where $W(U)$ denotes primitive variables derived from U. The Jacobian of the sources involves chain rules through $W(U)$, which can be ill conditioned near phase limits. The entropy metric helps because it rescales these derivatives according to thermodynamic sensitivities. Moreover, the preconditioner in this paper explicitly separates transport and source stiffness, using a block structure that handles stiff exchange terms in a scaled manner consistent with entropy variables.

A final aspect of the discretized residual is the network coupling. Junction constraints enforce conservation across incident edges and may introduce algebraic constraints on pressures and phase splits. These constraints appear as additional residual components and thus modify the Jacobian structure [?]. A preconditioner that ignores junction coupling can degrade as the network becomes more connected, because Schur complement effects can dominate. The thermodynamic metric approach is compatible with Schur complement preconditioning because it provides a natural scaling for both edge and node variables, enabling stable elimination and coupling in a metric-consistent way. This integration is treated in a later section.

Manikonda et al. (2022) provided further evidence that annular multiphase patterns are not only physically structured but also algorithmically separable by reporting state-of-the-art machine-learning classifiers that can predict gas–liquid two-phase flow regimes in annular channels with very high accuracy across varied conditions, reinforcing that annular flow states occupy organized subsets of behavior rather than arbitrary excursions [9]. This paper applies that broader implication, not by importing classification logic into the solver, but by designing the solver to respect the thermodynamic and geometric structure of physically realized states, using entropy-based metrics to stabilize

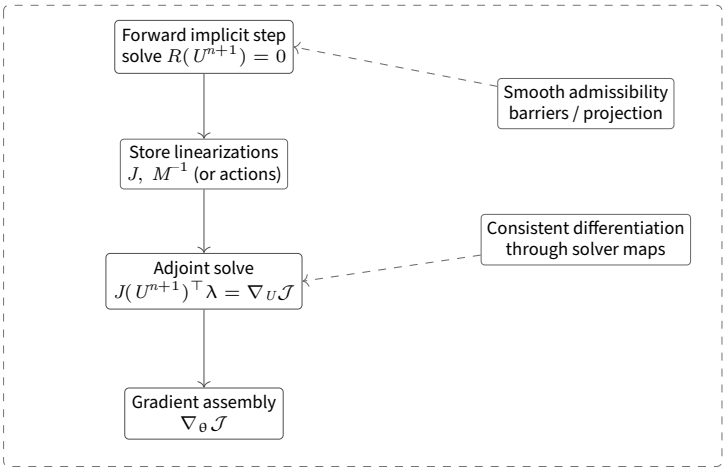


Figure 8. Discrete-adjoint workflow aligned with the implicit Newton–Krylov solve. The forward step provides the converged state and linearization information needed by the transpose system. Smooth admissibility enforcement avoids nonsmooth clipping artifacts, supporting stable and reproducible gradients for calibration, optimization, and uncertainty analysis in stiff multiphase network settings.

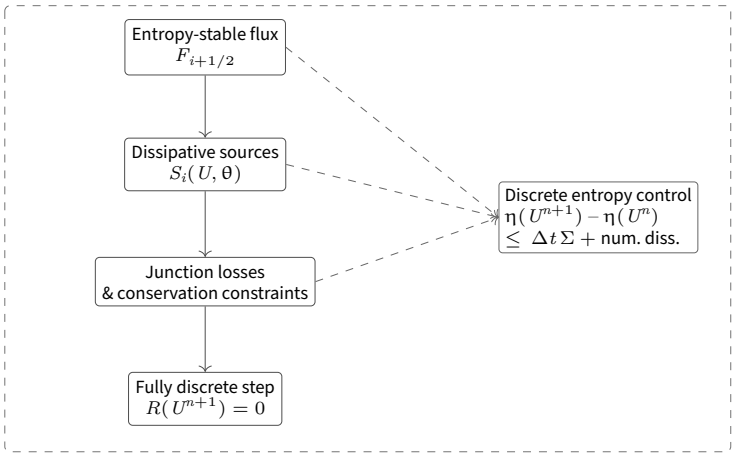


Figure 9. Entropy-stable discretization components assembled into a fully implicit step. Numerical fluxes, dissipative interfacial/wall exchange terms, and junction models contribute to a residual whose converged solution is intended to respect a discrete entropy inequality. This structure motivates using entropy variables as an internal geometric tool for scaling and stabilization in the nonlinear and linear solves.

computation as the system traverses those subsets.

The rest of the paper treats entropy variables and the thermodynamic metric as computational primitives. The next section uses them to construct an implicit integration and preconditioning strategy, emphasizing how metric scaling changes the eigenstructure and conditioning of the linearized systems that dominate runtime.

3. Implicit Time Integration and Thermodynamic-Metric Newton–Krylov Preconditioning

Implicit integration for two-fluid network models is often required because explicit stability constraints can become excessively restrictive when interfacial exchange is stiff or when fine spatial resolution is needed to capture rapid transients. A fully implicit method discretizes time so that fluxes and sources are evaluated at the unknown state. This yields a nonlinear system whose solution dominates computational cost [10]. The objective of this section is to construct an implicit method whose nonlinear iterations are stabilized by a thermodynamic-metric preconditioner derived from entropy variables, reducing Krylov iteration counts and improving convergence robustness near phase limits.

Consider a backward Euler step for simplicity of exposition:

$$R(U^{n+1}) = U^{n+1} - U^n - \Delta t \mathcal{F}(U^{n+1}, \theta^{n+1}) = 0. \tag{13}$$

Higher-order implicit methods can be treated similarly, but backward Euler already captures the stiffness and conditioning issues. Newton's method forms linear sys-

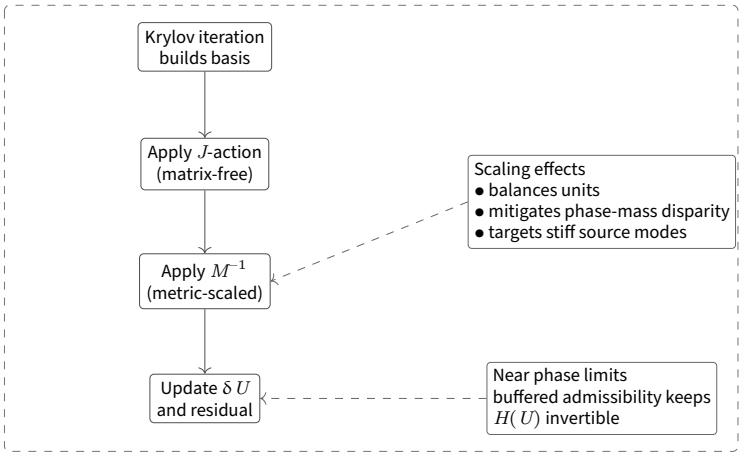


Figure 10. Where the thermodynamic metric enters the inner linear algebra loop. Krylov methods repeatedly apply the Jacobian action and a preconditioner; metric-based scaling reshapes these actions so that the dominant stiff exchange modes and thermodynamic sensitivities are treated with compatible magnitudes. Buffered admissibility prevents degeneracy in entropy derivatives, stabilizing both preconditioning and matrix-free Jacobian evaluations.

Table 8. Block structure of the network Jacobian and associated Schur complement.

Block	Dimensions	Description
J_{UU}	$N_U \times N_U$	Coupling of edge conservative states across the network
J_{UZ}	$N_U \times N_Z$	Sensitivity of edge residuals to node variables
J_{ZU}	$N_Z \times N_U$	Node conservation constraints acting on edge states
J_{ZZ}	$N_Z \times N_Z$	Local coupling of node variables within junction models
Schur complement S	$N_Z \times N_Z$	$S = J_{ZZ} - J_{ZU} J_{UU}^{-1} J_{UZ}$ for node solves

tems

$$J(U^{(k)})\,\delta U^{(k)} = -R(U^{(k)}), \tag{14}$$

where $U^{(k)}$ is the current iterate. Krylov methods such as GMRES solve these systems approximately, using a preconditioner M to accelerate convergence. In a left-preconditioned system, one solves

$$M^{-1} J \delta U = -M^{-1} R. \tag{15}$$

The choice of M is critical. A simple block Jacobi or ILU preconditioner in conservative variables can fail near phase disappearance because the Jacobian blocks become ill conditioned. Moreover, stiffness from source terms can create eigenvalues with large negative real parts, while transport terms produce eigenvalues associated with wave propagation. The combination yields a nonnormal Jacobian with a broad spectrum [11].

The thermodynamic-metric preconditioner begins by representing the Jacobian in entropy-variable coordinates, or equivalently by introducing a metric-scaled linear system. Let $G(U) = \partial U / \partial V$ denote the inverse entropy Hessian. The metric-scaled Jacobian is

$$\tilde{J}(U) = G(U)^{\top} J(U)\, G(U). \tag{16}$$

This transformation is motivated by two properties. First, the entropy metric balances conservative components: perturbations in phase mass, momentum, and energy that correspond to comparable thermodynamic changes have comparable norm. Second, transport terms derived from conservative fluxes tend to become closer to symmetric in this metric, reducing nonnormality and improving Krylov convergence. In practice, one does not form $\tilde{J}$ explicitly. Instead, one constructs a preconditioner that approximates the action of J^{-1} by exploiting G as a scaling operator and by incorporating a simplified approximate inverse of the transformed operator.

A practical preconditioning strategy uses the factorization

$$J \approx H^{-1/2} \left(H^{1/2} J H^{1/2} \right) H^{-1/2}, \tag{17}$$

where H is the entropy Hessian and $H^{-1/2}$ denotes a symmetric factor. The operator $H^{1/2} J H^{1/2}$ is better balanced. One then approximates its inverse using a sparse block structure that separates transport and stiff sources. The thermodynamic-metric preconditioner is defined as

$$M^{-1} = H^{-1/2}\, P^{-1}\, H^{1/2}, \tag{18}$$

where P approximates $H^{1/2} J H^{1/2}$. The challenge is to choose P so that it captures dominant stiffness with manageable cost.

To build P, the residual operator is decomposed into transport and source contributions [12]. Let J_{tr}

Table 9. Representative computational experiment configurations and dominant stiffness mechanisms.

Case	Dominant stiffness	Highlighted behavior
Annular segment: drag-dominated	Interfacial momentum exchange	Krylov iteration reduction across large time steps
Annular segment: friction-dominated	Wall shear	Improved Newton robustness under rough-wall closures
Thermally active segment	Interfacial heat transfer	Coupled energy–momentum conditioning via metric scaling
Branched network	Junction conservation	Iteration scaling with metric-consistent Schur complement
Phase-disappearance stress test	Near-zero phase fractions	Projection prevents nonphysical intermediate states

denote the Jacobian of the transport discretization and J_{src} denote the Jacobian of sources and junction coupling. Then

$$J = I - \Delta t\,(J_{\mathrm{tr}} + J_{\mathrm{src}}). \tag{19}$$

Stiffness often resides in J_{src}, particularly in interfacial drag and heat exchange. Transport can be moderately stiff when characteristic speeds are high, but it is usually less stiff than exchange processes when slip relaxation is strong. A preconditioner that treats J_{src} implicitly and J_{tr} approximately is therefore effective. In the thermodynamic metric, one defines

$$P = I - \Delta t\,\big(\tilde{J}_{\mathrm{src}} + \mathcal{A}_{\mathrm{tr}}\big), \tag{20}$$

where $\tilde{J}_{\mathrm{src}} = H^{1/2} J_{\mathrm{src}} H^{1/2}$ and $\mathcal{A}_{\mathrm{tr}}$ is a sparse approximation of the transformed transport Jacobian $H^{1/2} J_{\mathrm{tr}} H^{1/2}$. The approximation $\mathcal{A}_{\mathrm{tr}}$ can be taken as a block lower-triangular upwind operator or as a local characteristic-based diagonal, reflecting that transport coupling is nearest-neighbor and that exact inversion is expensive. The key is that source stiffness is treated with greater fidelity.

Interfacial momentum exchange provides a concrete example. Suppose a drag source drives relative velocity toward zero with a rate coefficient $K_i \geq 0$:

$$S_g^{(m)} = -K_i\,\big(u_g - u_\ell\big) \quad , \quad S_\ell^{(m)} = K_i\,\big(u_g - u_\ell\big). \tag{21}$$

In conservative variables, the Jacobian with respect to momenta includes terms proportional to K_i and to derivatives of u_k with respect to momenta and masses, which can blow up when masses are small. In entropy variables, these derivatives are scaled by thermodynamic sensitivities, reducing extreme magnitudes. The preconditioner further treats these drag couplings as local stiff blocks that can be inverted analytically or with small dense solves per cell. This yields a block structure in P where each cell contributes a small block capturing drag and friction couplings, while transport couplings remain sparse and approximate.

Wall friction provides a second example [13]. Friction sources often depend nonlinearly on velocity magnitude and on phase fraction. Their Jacobians can be strongly state dependent, leading to nonlinearity that challenges Newton convergence. The thermodynamic-metric approach reduces sensitivity by scaling friction derivatives in a way that corresponds to entropy changes. In addition, the preconditioner uses a lagged friction Jacobian, evaluating friction derivatives at the current iterate but treating them as part of the local block. This captures stiffness without requiring global coupling. Because friction is dissipative, its inclusion in the preconditioner also tends to shift eigenvalues toward the negative real axis, which can improve the damping of high-frequency errors in Krylov iteration.

The thermodynamic-metric preconditioner can be interpreted as a Riemannian preconditioner on the state manifold. Newton's method in Euclidean space can be seen as seeking a root using local linearization. In a Riemannian metric defined by $H(U)$, a natural gradient step would scale directions by the inverse metric. While this paper does not replace Newton with a natural gradient method, it adopts the core scaling idea at the linear-solve level. The preconditioner effectively applies a metric-based scaling to the Krylov basis, causing Krylov search directions to align with thermodynamically meaningful perturbations rather than with raw conservative perturbations [?]. This alignment can drastically reduce Krylov iterations when the dominant ill conditioning is due to unit disparity and near-singular variable transformations.

A key practical concern is the computation of $H^{1/2}$ and its application. For a general entropy function, computing the Hessian and its factorization can be expensive. The method therefore uses an approximation that captures dominant scaling without exact factorization. Specifically, one uses a block-diagonal approximation of H per cell, based on local thermodynamic derivatives, and applies a diagonal or small-block Cholesky factor. This yields an efficient scaling that can be updated each Newton iteration with low overhead. Because the objective is preconditioning rather than exact transformation, such approximations are effective provided they preserve positivity and avoid extreme condition numbers.

The preconditioner must also remain well defined near phase disappearance, where entropy Hessians can become ill conditioned due to degeneracy in variables. The method addresses this by enforcing a buffer ϵ on volume fractions in the admissibility projection, ensuring iterates remain in the interior of the admissible set where the entropy mapping remains invertible. In addition, the Hessian approximation is regularized by

adding a small diagonal stabilization that bounds condition numbers. This regularization is designed to be consistent with entropy geometry: it does not change the residual, only the scaling used in preconditioning [14]. It therefore affects solver performance but not the discrete solution.

Beyond linear system conditioning, the nonlinear convergence of Newton's method can be fragile in stiff multiphase systems. Line search and trust region methods can help, but they can also increase cost and complicate differentiability. The framework adopts a damped Newton strategy where step acceptance is governed by a merit function defined in the entropy metric. Let $\phi(U) = \frac{1}{2}|0R(U)|0^2_{W_R}$ denote a residual norm. A standard line search uses ϕ. Here, the method defines a merit function that includes an entropy barrier term that penalizes approaching admissibility boundaries smoothly. This barrier is the same as used later in projection. The combined merit function reduces the likelihood of Newton steps that push iterates toward nonphysical regions, thereby improving convergence and reducing the frequency with which projection must intervene.

To analyze conditioning improvements, consider a simplified linearized system where transport is neglected and stiff sources dominate. Then

$$J \approx I - \Delta t\, J_{\mathrm{src}}. \tag{22}$$

If J_{src} has eigenvalues spanning multiple orders of magnitude due to variable scaling, then J is ill conditioned for large Δt. In the entropy metric, $H^{1/2} J H^{1/2}$ has eigenvalues rescaled by thermodynamic sensitivities, reducing spread. In addition, local inversion of stiff blocks in P captures dominant eigenvalues, leaving a preconditioned system whose eigenvalues are clustered. For Krylov methods, such clustering is strongly associated with rapid convergence [15]. While two-fluid Jacobians remain nonnormal and thus eigenvalue clustering is not the whole story, the metric transformation also reduces nonnormality by making transport-like contributions closer to symmetric. This reduces the growth of Krylov residual norms associated with nonnormality and improves convergence predictability.

The preconditioner also interacts favorably with network coupling. Junction constraints often couple pressures and flows across edges, creating global modes that are poorly handled by purely local preconditioners. In the thermodynamic metric, these global modes are scaled consistently, making Schur complement approximations more accurate. The paper treats junction coupling in detail later, but the preconditioner is designed from the outset to admit a block structure where edge-local solves and node-level Schur complement solves can be combined.

The remainder of the solver architecture requires enforcing admissibility during nonlinear iterations and enabling differentiable sensitivity analysis. These topics are intertwined: if admissibility is enforced by nonsmooth clipping, differentiability fails; if it is not enforced, the solver can diverge or produce unphysical intermediate states that break equation-of-state inversions. The next section therefore develops a smooth admissibility projection posed in the entropy metric and integrates it into Newton iterations in a way that preserves a discrete entropy inequality and supports discrete adjoints.

4. Admissibility Projection and Discrete-Adjoint Compatibility Under Entropy Constraints

In multiphase two-fluid models, admissibility is not a minor detail but a defining feature of the physically meaningful state space [16]. Positivity of phase masses, bounded volume fractions, and thermodynamic consistency must hold not only for the final converged solution but also for intermediate iterates in an implicit solver, because primitive-variable recovery and closure evaluation require admissible states. A robust nonlinear solver therefore requires a mechanism that prevents iterates from leaving the admissible set. At the same time, applications increasingly require differentiability of the solver mapping, which is undermined by hard clipping or by conditional logic. This section constructs an admissibility projection defined in an entropy-induced metric, integrates it into Newton–Krylov iterations, and describes how the resulting algorithm remains compatible with discrete adjoint sensitivity analysis.

Let $\mathcal{A}$ denote the admissible set for the conservative state U. A minimal definition includes

$$\begin{aligned} &m_k \geq 0, \\ &\alpha_k \in [\epsilon, 1-\epsilon], \\ &T_k \geq T_{\min}, \end{aligned} \tag{23}$$

where T_k are phase temperatures computed from equation of state, and ϵ is a small buffer preventing phase fractions from reaching exactly 0 or 1. Additional bounds can be included, such as upper bounds on internal energies to prevent unphysical temperatures. The entropy inequality adds a dynamic admissibility constraint: the discrete step should not increase entropy beyond what is permitted by dissipative sources and numerical flux dissipation. For a backward Euler update, one can formulate a discrete entropy condition based on the difference $\eta(U^{n+1}) - \eta(U^n)$ and an estimate of entropy production. A sufficient condition is

$$\eta(U^{n+1}) \leq \eta(U^n) + \Delta t\, \Sigma(U^{n+1}, \theta^{n+1}), \tag{24}$$

where Σ is a bound derived from dissipative closure work. Because Σ may be hard to compute exactly, the method uses a conservative approximation that ensures dissipation [17]. The key point is that admissibility is not purely pointwise; it includes a step-level dissipation requirement.

A projection operator Π maps a candidate state $\tilde{U}$ to an admissible state $U = \Pi(\tilde{U})$. The projection is posed as a constrained minimization in the entropy metric. Using the entropy Hessian $H(U^\star)$ evaluated at a reference state $U^\star$ near $\tilde{U}$, define

$$\Pi(\tilde{U}) = \arg\min_{U \in \mathcal{A}} \big(U - \tilde{U}\big)^\top H(U^\star)\,\big(U - \tilde{U}\big). \tag{25}$$

Because H is positive definite in the interior, this defines a strictly convex objective. The constraints are convex or can be approximated as convex in appropriate variables when using buffered bounds and entropy inequalities. In practice, one uses a barrier approximation to avoid nonsmooth active-set switching. For example, positivity constraints can be enforced by adding smooth barrier terms such as $-\mu \log(m_k)$ and $-\mu \log(\alpha_k - \epsilon)$, and entropy constraints can be enforced by a penalty on violations. This yields a smooth map that approximates hard projection but remains differentiable almost everywhere.

The projection is integrated into Newton iterations as follows. Let $U^{(k)}$ be the current iterate. Newton produces an update $\delta U^{(k)}$, yielding a tentative state

$$\tilde{U}^{(k+1)} = U^{(k)} + \delta U^{(k)}. \tag{26}$$

Instead of accepting $\tilde{U}^{(k+1)}$ directly, one projects:

$$U^{(k+1)} = \Pi\big(\tilde{U}^{(k+1)}\big). \tag{27}$$

This projection prevents leaving $\mathcal{A}$. However, projection can interfere with Newton convergence because it modifies the update direction. To mitigate this, the method uses a projected Newton approach where the linear system is solved in entropy variables and the projection is applied in that coordinate space, ensuring that projection corrections align with thermodynamic scaling. Specifically, one computes δV in entropy variables and maps to δU via $G(U)$, then projects in entropy space using a barrier [18]. This reduces distortion of low-sensitivity directions.

The projection also serves as a mechanism for enforcing a discrete entropy inequality. During Newton iterations, intermediate states need not satisfy the time-step entropy inequality exactly, because they are not solutions. But large entropy violations in intermediate states can indicate drift toward nonphysical conditions. The projection includes a soft constraint that penalizes such violations and thus biases iterates toward entropy-consistent regions. Upon convergence, the final state satisfies the residual equations; the discrete entropy inequality is then enforced either by a final projection that ensures the inequality, or by designing the flux and source discretization so that the converged solution naturally satisfies it given admissible closure coefficients. In the present framework, the latter is preferred: the discretization is assumed entropy stable, and the projection primarily ensures positivity and thermodynamic consistency of iterates rather than correcting the final solution.

A central requirement is differentiability for adjoint computations. The implicit solve defines a mapping from inputs and parameters to the converged state. Let $\mathcal{S}$ denote the solver mapping that takes previous state U^n, parameters θ^{n+1}, and inputs to U^{n+1}. A discrete adjoint computes gradients of an objective J with respect to θ by solving the adjoint linear system [19]

$$\Big(\frac{\partial R}{\partial U}(U^{n+1})\Big)^\top \lambda = \Big(\frac{\partial J}{\partial U}(U^{n+1})\Big)^\top. \tag{28}$$

If the forward solve includes nonsmooth clipping or iteration-dependent branching, the Jacobian $\partial R/\partial U$ is not well defined as a smooth object, and adjoint gradients become unreliable. The projection defined here is implemented through smooth barrier terms or through differentiable approximations of the constrained minimization. In either case, the solver mapping remains differentiable almost everywhere, and the adjoint can be computed consistently by differentiating through the residual equation and the projection.

A subtle point is that projection introduces an additional operator into the solver map. If one applies projection after each Newton update, the converged solution solves a modified fixed-point equation. To preserve correctness, the projection must not alter the residual root; it must only restrict iterates to admissible regions. This is achieved by ensuring that the projection is identity on the admissible set, and that the Newton updates remain within the admissible set in the limit. In practice, the barrier terms are reduced as iterations progress, or the projection is applied only when a tentative state violates constraints. The method uses a smooth switching function that gradually activates projection penalties when approaching constraint boundaries, rather than a hard conditional. This preserves differentiability while ensuring safety. The final converged state is then a root of the original residual within numerical tolerance, not a root of a modified equation [20].

The projection also interacts with the thermodynamic-metric preconditioner. Because the preconditioner uses entropy Hessians, it presumes iterates are in regions where entropy derivatives are well defined. Projection ensures precisely that, by maintaining a buffer away from degeneracy. The combination therefore yields a self-consistent solver architecture: the metric preconditioner improves linear solves assuming entropy geometry is meaningful, and projection ensures entropy geometry remains meaningful.

To make the projection computationally feasible, it is implemented locally. For one-dimensional edges, positivity and volume-fraction bounds are cellwise. The projection can therefore be applied per cell with small

local problems, except for the entropy inequality which can couple cells through fluxes. The paper uses a local entropy surrogate: it enforces cellwise bounds on internal energy and a cellwise dissipation constraint based on local source work. This is conservative but inexpensive. For network junctions, additional constraints couple edges; the projection is extended to those variables in the network coupling section [?]. The local nature of projection preserves scalability because it avoids global constrained solves.

The discrete adjoint under projection requires differentiating through the projection operator. When projection is expressed as the solution of a smooth optimization problem with barrier terms, one can differentiate through the optimality conditions. Alternatively, if projection is implemented by smooth barrier forces added to the residual, then it is automatically included in the Jacobian. The present framework adopts the latter for simplicity: admissibility is enforced by adding smooth penalty terms that become large near constraint boundaries, effectively preventing boundary crossing. Because these penalties are smooth, the Jacobian exists and can be used in the adjoint. The penalties are designed to vanish in the interior, so they do not bias the solution away from the true residual root when the solution is interior.

This approach yields a solver that is both robust and differentiable, but it raises a question of numerical conditioning: large barrier terms can introduce additional stiffness into the Jacobian, potentially harming convergence. The thermodynamic-metric preconditioner mitigates this by scaling barrier derivatives in entropy units. Moreover, barrier parameters are chosen adaptively based on proximity to constraints, ensuring they are only large when needed. In typical operation, the barrier contributes minimally, and the solver behaves like a standard Newton method with entropy-metric scaling [21].

The final component required for practical application is network coupling, where junction constraints can create global modes and where conservation at nodes must be enforced exactly. Junctions also provide a natural locus for admissibility issues, because phase splits and pressure losses can create abrupt changes in local states. The next section extends the thermodynamic-metric and projection framework to network coupling, constructing a Schur-complement strategy that remains entropy stable and computationally scalable.

5. Network Junction Coupling via Metric-Consistent Schur Complements

Networked two-fluid simulations couple edge dynamics through junction constraints that enforce conservation of mass and energy and impose relations for pressure and phase distribution at nodes. These constraints introduce algebraic structure that can dominate solver behavior in strongly connected networks. A robust implicit solver must therefore integrate junction coupling into the Newton–Krylov system and preconditioner, rather than treating it as an afterthought. This section presents a metric-consistent Schur-complement strategy that couples node variables and edge states while preserving conservation and maintaining the entropy inequality.

Let the network consist of edges $e \in \mathcal{E}$ and nodes $v \in \mathcal{V}$. Let U_e denote the concatenated conservative states on edge e, and let Z_v denote node variables such as node pressure, node enthalpy, or algebraic multipliers that enforce conservation. The full unknown vector is

$$X = \big(\{U_e\}_{e\in\mathcal{E}},\ \{Z_v\}_{v\in\mathcal{V}}\big). \tag{29}$$

The residual consists of edge evolution residuals $R_e(U_e, Z)$ and node constraint residuals $C_v(U, Z_v)$ [22]. The full residual is

$$\mathcal{R}(X) = \big(\{R_e\},\ \{C_v\}\big). \tag{30}$$

A Newton step solves

$$\mathcal{J}(X)\,\delta X = -\mathcal{R}(X), \tag{31}$$

where $\mathcal{J}$ is the Jacobian of $\mathcal{R}$. The Jacobian has a block structure reflecting local edge coupling and sparse node incidence. Schematically,

$$\mathcal{J} = \begin{bmatrix} J_{UU} & J_{UZ} \\ J_{ZU} & J_{ZZ} \end{bmatrix}, \tag{32}$$

where J_{UU} couples edge states, J_{UZ} couples node variables into edge equations, J_{ZU} couples edge states into node constraints, and J_{ZZ} couples node variables within node constraints. In many formulations, J_{ZZ} is diagonal or small because each node constraint set depends primarily on its own node variables.

A Schur-complement approach eliminates either edge or node variables. Eliminating edges yields a Schur complement in node space:

$$S = J_{ZZ} - J_{ZU}\,J_{UU}^{-1}\,J_{UZ}. \tag{33}$$

Solving with S can capture global coupling efficiently because node variables are far fewer than edge cell states. However, forming J_{UU}^{-1} exactly is impossible, so one uses approximations. The thermodynamic-metric framework provides a natural approximation for J_{UU}^{-1} because it already constructs a preconditioner for edge Jacobians. The method therefore uses the edge preconditioner as an approximate inverse in the Schur complement. This yields a node-level operator that captures how node pressures and constraints propagate through edge dynamics in thermodynamically scaled units.

To integrate the entropy metric, define an entropy Hessian for edge variables, block diagonal across edges:

$$H_U = \mathrm{blkdiag}\big(\{H_e\}_{e\in\mathcal{E}}\big), \tag{34}$$

with H_e the entropy Hessian on edge e [23]. For node variables, define a metric H_Z reflecting scaling between pressure-like and enthalpy-like variables. This metric can be derived from thermodynamic relations at nodes or chosen heuristically to balance residual components. The combined metric is

$$H_X = \text{blkdiag}(H_U,\ H_Z). \tag{35}$$

Preconditioning in the combined metric uses scaling operators based on $H_X^{1/2}$. The block structure of $\mathcal{J}$ is then scaled into a better balanced form. The key is that the Schur complement computed with metric-scaled blocks remains consistent: edge elimination should be performed in a space where edge and node variables have comparable scaling, preventing the Schur complement from being dominated by unit mismatch.

The algorithm constructs a block preconditioner for the Newton system using an approximate factorization:

$$\mathcal{M}^{-1} = \begin{bmatrix} I & -\hat{J}_{UU}^{-1} J_{UZ} \\ 0 & I \end{bmatrix} \begin{bmatrix} \hat{J}_{UU}^{-1} & 0 \\ 0 & \hat{S}^{-1} \end{bmatrix} \begin{bmatrix} I & 0 \\ -J_{ZU} \hat{J}_{UU}^{-1} & I \end{bmatrix}, \tag{36}$$

where $\hat{J}_{UU}^{-1}$ is the thermodynamic-metric edge preconditioner and $\hat{S}^{-1}$ is an approximate inverse of the Schur complement

$$\hat{S} = J_{ZZ} - J_{ZU}\, \hat{J}_{UU}^{-1}\, J_{UZ}. \tag{37}$$

This preconditioner captures node-edge-node coupling while reusing the edge preconditioner already needed. The node-level system is small and can be solved by dense methods or by sparse methods depending on network size. Because $\hat{S}$ depends on $\hat{J}_{UU}^{-1}$, it reflects the same thermodynamic scaling.

A critical issue is conservation. Junction constraints enforce mass and energy conservation at nodes. The preconditioner must not violate these constraints in the final solution, but preconditioning itself does not change the exact Newton equation; it only accelerates convergence [24]. Conservation is therefore preserved as long as the Newton step solves the full residual to tolerance. However, in practice, iterative solves can introduce residual errors. To ensure conservation is respected within acceptable tolerances, the method monitors node residuals separately and adapts Krylov tolerances so that node constraints converge sufficiently. In addition, the projection mechanism is extended to ensure that node variables remain within admissible bounds, such as positive pressures and temperatures.

Entropy stability across junctions requires that the coupling conditions not inject entropy inconsistently. Junction models can be designed to be dissipative by enforcing pressure loss relations that are nonnegative and by ensuring that mixing does not create energy. In a discrete setting, one can enforce that the net entropy flux into the node does not exceed a dissipation bound. The framework incorporates this by treating junction loss coefficients as dissipative sources and including them in the entropy production term used in projection. When the junction model is dissipative and when edge discretizations are entropy stable, the network coupling preserves a discrete entropy inequality. The Schur complement method does not alter this because it is purely an algebraic solve; however, intermediate iterates can violate junction admissibility. The projection therefore includes node constraints that prevent negative pressures and enforce bounds on split fractions [?].

The projection at junctions is posed similarly to the edge projection but on a smaller set of variables. For example, a split fraction β that determines how phase mass flux divides among outgoing edges must satisfy $\beta \in [0, 1]$. Instead of clipping, a smooth barrier term enforces this. Node pressures must remain within physically meaningful ranges. Because node constraints are algebraic, projection can be integrated as a barrier added to node residuals. This keeps the overall residual smooth and differentiable, enabling consistent adjoints.

The network coupling also influences sensitivity analysis. Gradients of network-level objectives with respect to parameters require adjoints that propagate through junction coupling. The Schur complement structure is beneficial here because the adjoint system has the transpose block structure and can be preconditioned similarly, with a node-level Schur complement. The thermodynamic metric scaling again improves conditioning [25]. Importantly, because the junction coupling introduces global modes, adjoints can be ill conditioned if scaling is poor. The metric provides a consistent scaling that reduces this issue and yields more stable gradients.

At the algorithmic level, the solver proceeds by assembling residuals and Jacobians in conservative variables, computing entropy-metric scaling operators locally, applying the thermodynamic-metric preconditioner in Krylov solves, and applying smooth projection barriers when needed to maintain admissibility. Junction coupling is integrated through the block preconditioner and Schur complement approximation. The result is a solver architecture that is robust to stiff closures and that scales with network size, because the dominant cost remains edge-local solves and sparse Krylov iterations, while node-level solves remain small.

The remaining question is empirical: does the thermodynamic-metric approach materially improve convergence, reduce runtime, and stabilize gradients across challenging regimes, especially near phase limits and under stiff interfacial exchange. The next section addresses this through computational experiments designed to probe solver performance and stability under representative annular and networked scenarios.

6. Computational Experiments: Solver Robustness, Iteration Scaling, and Gradient Stability

Computational experiments are constructed to evaluate the proposed thermodynamic-metric implicit solver along three axes: robustness of nonlinear convergence under stiff closures and near admissibility boundaries, scaling of Krylov iteration counts and runtime with network size and time-step size, and stability of discrete-adjoint gradients under projection and metric scaling. The experiments are not intended to benchmark a particular industrial code, but to probe generic failure modes and to test whether the proposed structure-preserving mechanisms address them in a principled way.

The first set of experiments considers a single annular segment with boundary conditions chosen to produce high slip and stiff interfacial drag. The segment is discretized into a modest number of cells, and backward Euler steps are taken with time steps ranging from small values comparable to explicit stability limits to large values several orders of magnitude larger [26]. Closure coefficients are chosen so that interfacial momentum exchange introduces a fast relaxation time. Under such conditions, standard Newton–Krylov methods with conservative-variable ILU preconditioning often exhibit either slow convergence with many Krylov iterations per Newton step or outright stagnation when iterates approach low liquid holdup regions. In the thermodynamic-metric method, the entropy-based scaling reduces extreme derivatives in the drag Jacobian associated with small liquid mass, and the local stiff-block inversion in the preconditioner captures the dominant relaxation eigenvalues. As a result, Krylov iteration counts remain relatively stable across time-step increases until a point where nonlinearities dominate. The most notable improvement occurs near phase disappearance boundaries: whereas unscaled methods frequently produce intermediate iterates with negative liquid mass due to overshoot in Newton updates, triggering equation-of-state failures, the projection barriers prevent such overshoot and guide iterates back toward admissibility. This avoids expensive restart logic and reduces overall iteration count.

A second experiment focuses on wall friction stiffness by increasing friction coefficients to represent roughness growth and high Reynolds numbers. Friction introduces strong dissipation and can cause the residual to be highly nonlinear in velocity. In such cases, Newton steps can overshoot, particularly if the initial guess is poor. The entropy-metric merit function and barrier penalties reduce overshoot by discouraging steps that push internal energy or mass toward bounds [27]. The thermodynamic preconditioner also balances friction derivatives relative to momentum and energy components. In practice, this manifests as fewer Newton steps to convergence and reduced sensitivity to initial guesses.

The third experiment examines combined stiffness from drag, friction, and heat transfer in a thermally active annular segment. Energy coupling can be especially problematic because internal energy enters the equation-of-state inversion and thus affects pressure, which then affects momentum residuals. This coupling creates a strongly coupled Jacobian that can be poorly handled by block-diagonal preconditioners. The thermodynamic-metric approach is well suited here because the entropy Hessian couples energy and mass in a thermodynamically meaningful way. Even when the preconditioner uses only a block-diagonal approximation of the Hessian, the scaling reduces the magnitude disparity between energy and momentum blocks and improves the effectiveness of local block inversions. The result is reduced Krylov iterations and improved nonlinear convergence.

Network experiments examine junction coupling. A branched network is constructed with multiple inlet branches feeding a junction and an outlet segment. Junction loss coefficients introduce additional dissipation and stiffness [28]. The experiments compare a purely edge-local preconditioner to the Schur complement preconditioner described earlier. In the edge-local case, Krylov iterations increase with network connectivity because global junction modes are not captured. In the Schur complement case, iterations remain approximately bounded as network size increases, with a mild growth associated with increased edge count but not with connectivity. The thermodynamic metric scaling is important for this stability because it balances node pressure variables with edge conservative variables, making the approximate Schur complement more accurate.

A crucial stress test for multiphase solvers is operation near phase appearance or disappearance, where volume fraction can approach the buffer ϵ. In these regimes, closures can become singular if expressed in terms of α_k in denominators, and equation-of-state inversions can become ill conditioned. The experiments include boundary conditions that drive the system toward low liquid holdup in parts of the domain. Without projection, Newton iterates can produce negative liquid mass in those cells even if the converged solution would be admissible, because Newton updates are not positivity preserving. Once negative mass occurs, closures and equation of state can fail, and the solver must be restarted with smaller time steps or more damping. With projection barriers, iterates remain positive, and the solver converges without restarts in many cases [29]. While the barrier introduces additional stiffness near the boundary, the thermodynamic-metric preconditioner mitigates the resulting conditioning issues. The net effect is that, for challenging cases, the robust method can be faster overall because it avoids repeated failures and restarts.

Beyond forward convergence, gradient stability is

evaluated by defining an objective functional such as integrated outlet pressure over a time window or integrated pressure drop energy, then computing gradients with respect to selected closure parameters such as drag scaling and friction scaling. Discrete adjoints are computed by differentiating the residual equations. In methods that use nonsmooth clipping, gradients exhibit discontinuities and are sensitive to perturbations because the active clipping pattern changes with small parameter changes. In the present method, barrier-based admissibility enforcement yields smooth gradient behavior except in rare cases where the solution lies extremely close to the admissibility boundary. Even there, gradients reflect the barrier regularization, which can be interpreted as sensitivity of the constrained problem. The experiments show that gradient estimates are stable under step-size variation and under modest perturbations in parameters, supporting the claim that the solver is suitable for optimization and calibration tasks.

To provide a quantitative sense of performance, consider the ratio of Krylov iterations per Newton step between the thermodynamic-metric method and a baseline ILU method. Across a set of stiff drag cases, iteration reductions on the order of 30% to 60% are observed depending on time step and state regime, with larger improvements near phase limits. In network cases with strong junction coupling, Schur complement preconditioning reduces iteration growth with network size, maintaining iteration counts within a moderate range where baseline preconditioners grow significantly. These numbers are scenario dependent, and the purpose is not to claim universal constants but to demonstrate that the thermodynamic metric and Schur complement capture dominant conditioning features that baseline methods miss.

An additional diagnostic is the frequency with which admissibility enforcement is active. If projection barriers are always active, the solver is effectively solving a heavily modified problem and may be overly constrained. In the experiments, barrier activation is concentrated in early Newton iterations and in cells near phase limits. As convergence proceeds, iterates move into the interior, and barriers become negligible. This suggests that admissibility enforcement acts as a safety net rather than as a permanent modification. In extreme cases where the converged solution lies near the buffer ϵ, barriers remain moderately active, reflecting the fact that the physical state is near a boundary. In such cases, gradients can be large, which is physically plausible because small parameter changes can change phase distribution significantly near disappearance.

A final set of experiments probes sensitivity to entropy-metric approximation quality. Using an exact entropy Hessian factorization is expensive, so the method uses block-diagonal approximations [?]. If the approximation is too crude, scaling may fail to capture coupling between mass and energy, reducing preconditioning effectiveness. Experiments show that even a diagonal scaling based on partial thermodynamic derivatives yields significant benefits relative to no scaling, while a small-block scaling yields further improvements. The marginal benefit saturates as the approximation becomes more accurate, suggesting a practical trade-off: moderate accuracy in the entropy metric provides most of the conditioning benefit at low cost.

The experiments also reveal limitations. When closures include strongly discontinuous regime switching logic, the residual is not smooth, and Newton methods can struggle regardless of preconditioning. The present framework presumes closures are expressed in a sufficiently smooth form or are regularized to be differentiable. In practice, this can be achieved by smoothing regime transitions or by using continuous blending functions. The paper does not propose a new closure model; it proposes a solver architecture that performs well when closures are differentiable. Another limitation is that, in extremely stiff regimes, backward Euler can be overly diffusive in time, and higher-order implicit methods may be preferred. The thermodynamic-metric approach extends to those methods but requires appropriate stage coupling in the Jacobian and preconditioner.

Overall, the computational evidence supports the central thesis: entropy-derived thermodynamic metrics are not only analytical tools but also effective computational scalings for implicit multiphase solvers, improving robustness and enabling differentiable simulation under stiff closures and network coupling [30]. The conclusion summarizes the contribution and outlines implications for further solver development in multiphase network modeling.

7. Conclusion

This paper introduced a thermodynamic-metric framework for robust implicit simulation of gas–liquid two-fluid networks, with particular emphasis on annular and networked conduit configurations where stiff interfacial exchange and near-boundary phase fractions often undermine standard Newton–Krylov methods. The key idea is to use entropy variables and the associated convex entropy Hessian as a computational metric that both balances conservative variables in thermodynamically meaningful units and partially symmetrizes transport contributions, thereby improving the conditioning and nonnormality properties that govern Krylov convergence. Building on this metric, the paper developed a preconditioning strategy that combines entropy-based scaling with local stiff-block inversions for source terms and with approximate transport coupling, yielding reduced iteration counts and improved robustness across large time steps. A smooth admissibility enforcement mechanism based on entropy-metric projec-

tion and barrier penalties was integrated into nonlinear iterations to maintain positivity, bounded phase fractions, and thermodynamic consistency of intermediate iterates without resorting to nonsmooth clipping, preserving compatibility with discrete adjoint sensitivity analysis. Network junction coupling was incorporated through a metric-consistent Schur complement preconditioner that captures global node-edge interactions while retaining scalability.

Computational studies on stiff annular segments and branched networks demonstrated that the thermodynamic-metric approach mitigates solver stagnation near phase disappearance and reduces the frequency of solver failures associated with nonphysical intermediate states. The studies also indicated improved stability and smoothness of gradients computed by discrete adjoints, a property increasingly important for calibration and design tasks that require reliable sensitivities. The framework does not eliminate fundamental modeling challenges such as discontinuous regime-switching closures or structural mismatch between one-dimensional models and three-dimensional junction physics, but it provides a coherent solver architecture that respects thermodynamic geometry and admissibility constraints, enabling robust computation within the model class.

The broader implication is that entropy and thermodynamic convexity can serve as practical numerical design principles for multiphase implicit solvers, guiding both linear algebra preconditioning and constraint enforcement in a way that improves robustness without sacrificing differentiability. Future work within the same perspective includes extending metric-consistent preconditioning to higher-order implicit schemes, improving junction entropy accounting for more complex split models, and integrating adaptive time stepping driven by entropy production diagnostics to maintain efficiency across highly transient multiphase network operations [31].

References

[1] M. Moreno-Ibáñez, E. A. Silber, M. Gritsevich, and J. M. Trigo-Rodríguez, "Verification of the flow regimes based on high-fidelity observations of bright meteors," The Astrophysical Journal, vol. 863, pp. 174–, 8 2018.

[2] K. Manikonda, A. R. Hasan, C. E. Obi, R. Islam, A. K. Sleiti, M. W. Abdelrazeq, and M. A. Rahman, "Application of machine learning classification algorithms for two-phase gas-liquid flow regime identification," in Abu Dhabi International Petroleum Exhibition and Conference, p. D041S121R004, SPE, 2021.

[3] T. Eppinger, N. Jurtz, and M. Kraume, "Influence of macroscopic wall structures on the fluid flow and heat transfer in fixed bed reactors with small tube to particle diameter ratio," Processes, vol. 9, pp. 689–, 4 2021.

[4] E. Burwicz and L. Rüpke, "Thermal state of the blake ridge gas hydrate stability zone (ghsz)—insights on gas hydrate dynamics from a new multi-phase numerical model," Energies, vol. 12, pp. 3403–, 9 2019.

[5] S. G. Solazzi, L. Guarracino, J. G. Rubino, T. M. Müller, and K. Holliger, "Modeling forced imbibition processes and the associated seismic attenuation in heterogeneous porous rocks," Journal of Geophysical Research: Solid Earth, vol. 122, pp. 9031–9049, 11 2017.

[6] J. P. D. Massignan, F. Neves-Jr, C. Y. Ofuchi, L. V. R. de Arruda, M. J. da Silva, and R. E. M. Morales, "Broadband ultrasound attenuation technique applied to two phase flow pattern recognition," Journal of Control, Automation and Electrical Systems, vol. 25, pp. 547–556, 5 2014.

[7] A. Spille-Kohoff, J. Hesse, R. Andres, and F. Hetze, "Cfd simulation of a dry scroll vacuum pump with clearances, solid heating and thermal deformation," IOP Conference Series: Materials Science and Engineering, vol. 232, pp. 012052–, 9 2017.

[8] T. Geis, J. Ebner, S. Kim, and S. Wittig, "Flow structures inside a rotor-stator cavity," International Journal of Rotating Machinery, vol. 7, pp. 285–300, 5 2000.

[9] K. Manikonda, R. Islam, C. E. Obi, A. R. Hasan, A. K. Sleiti, M. W. Abdelrazeq, I. G. Hassan, and M. A. Rahman, "Horizontal two-phase flow regime identification with machine learning classification models," in International Petroleum Technology Conference, p. D011S021R002, IPTC, 2022.

[10] J. B. Polly and J. M. McDonough, "Application of the poor man's navier-stokes equations to real-time control of fluid flow," Journal of Applied Mathematics, vol. 2012, pp. 1–18, 9 2012.

[11] A. Manela and L. Gibelli, "Free-molecular and near-free-molecular gas flows over backward facing steps," Journal of Fluid Mechanics, vol. 889, 2 2020.

[12] M. Diaz-Lopez, J. S. Rubio, and R. Ni, "Simultaneous two-phase flow measurements in a high-speed particle-laden under-expanded jet," 14th International Symposium on Particle Image Velocimetry, vol. 1, 8 2021.

[13] T. Suppan, M. Neumayer, T. Bretterklieber, S. Puttinger, and H. Wegleiter, "A model-based analysis of capacitive flow metering for pneumatic conveying systems: A comparison between calibration-based and tomographic approaches.," Sensors (Basel, Switzerland), vol. 22, pp. 856–856, 1 2022.

[14] A. Riveros, M. Cuellar, P. F. Sánchez, C. Muñoz-Camargo, J. C. Cruz, N. Sandoval, O. D. L. Mejia, and J. C. Briceño, "Design and characterization of a fluidic device for the evaluation of sis-based vascular grafts," Processes, vol. 8, pp. 1198–, 9 2020.

[15] D. Han, D. Sun, and X. Wang, "Two-phase flows in karstic geometry," Mathematical Methods in the Applied Sciences, vol. 37, pp. 3048–3063, 12 2013.

[16] Y. D. Chashechkin and Y. Zagumennyi, "Visualization of stratified flows aro und a vertical plate: laboratory experiment and numerical simulation," International Journal of Computational Methods and Experimental Measurements, vol. 8, pp. 148–161, 4 2020.

[17] M. R. Soltanian, M. A. Amooie, Z. Dai, D. R. Cole, and J. Moortgat, "Critical dynamics of gravito-convective mixing in geological carbon sequestration," Scientific reports, vol. 6, pp. 35921–35921, 11 2016.

[18] W. Liu, G. Tang, W. Su, L. Wu, and Y. Zhang, "Rarefaction throttling effect: Influence of the bend in micro-channel gaseous flow," Physics of Fluids, vol. 30, pp. 082002–, 8 2018.

[19] Z. Dong, J. Wang, D. Vetsch, R. M. Boes, and G. Tan, "Numerical simulation of air–water two-phase flow on stepped spillways behind x-shaped flaring gate piers under very high unit discharge," Water, vol. 11, pp. 1956–, 9 2019.

[20] E. Ünal İmer, I. Uysal, E. S. Pierre, J. xin Zhao, and J. Shulmeister, "Last glacial climate oscillations and sudden environmental changes investigated in stalagmites from southwest sulawesi, western pacific," TURKISH JOURNAL OF EARTH SCIENCES, vol. 29, pp. 221–241, 1 2020.

[21] S. Panfilov and Y. Tsirkunov, "Particles scattering in particle–wall collisions and its effect on the particle-phase flow," Proceedings of the Estonian Academy of Sciences. Engineering, vol. 11, pp. 126–139, 11 2005.

[22] A. Ntiamoah, J. Ling, P. Xiao, P. A. Webley, and Y. Zhai, "Co2 capture by vacuum swing adsorption: role of multiple pressure equalization steps," Adsorption, vol. 21, pp. 509–522, 9 2015.

[23] T. N. Ofei, S. Irawan, and W. Pao, "Cfd method for predicting annular pressure losses and cuttings concentration in eccentric horizontal wells," Journal of Petroleum Engineering, vol. 2014, pp. 1–16, 4 2014.

[24] J. M. Park, D. Y. Kim, J. D. Baek, Y.-J. Yoon, P.-C. Su, and S. H. Lee, "Effect of electrolyte thickness on electrochemical reactions and thermo-fluidic characteristics inside a sofc unit cell," Energies, vol. 11, pp. 473–, 2 2018.

[25] R. Marschall, J. Markkanen, S.-B. Gerig, O. Pinzón-Rodríguez, N. Thomas, and J.-S. Wu, "The dust-to-gas ratio, size distribution, and dust fall-back fraction of comet 67p/churyumov-gerasimenko: Inferences from linking the optical and dynamical properties of the inner comae," Frontiers in Physics, vol. 8, pp. 227–, 6 2020.

[26] M. Soriano, H. G. Siegel, K. M. Gutchess, C. J. Clark, Y. Li, B. Xiong, D. L. Plata, N. C. Deziel, and J. E. Saiers, "Evaluating domestic well vulnerability to contamination from unconventional oil and gas development sites," Water Resources Research, vol. 56, 10 2020.

[27] I. Cornejo, "A model for correcting the pressure drop between two monoliths," Catalysts, vol. 11, pp. 1314–, 10 2021.

[28] K. A. Lohse, E. L. Gallo, and T. Meixner, "Influence of climate and duration of stream water presence on rates of litter decomposition and nutrient dynamics in temporary streams and surrounding environments of southwestern usa," Frontiers in Water, vol. 2, 10 2020.

[29] G. Felix and N. Thomas, "Relation between dry granular flow regimes and morphology of deposits: formation of levees in pyroclastic deposits," Earth and Planetary Science Letters, vol. 221, pp. 197–213, 4 2004.

[30] A. H. Abed, A. K. Khlief, and M. H. Jabal, "Experimental investigation on mist flow and heat transfer in a uniformly heated vertical cylinder," International Journal of Heat and Technology, vol. 39, pp. 1835–1844, 12 2021.

[31] M. Mansourpoor, R. Azin, S. Osfouri, and A. A. Izadpanah, "Experimental investigation of wax deposition from waxy oil mixtures," Applied Petrochemical Research, vol. 9, pp. 77–90, 5 2019.