International Journal of Academic and Applied Research (IJAAR)
ISSN: 2643-9603
Vol. 10 Issue 10, October - 2026, Pages: 49-63

A Phase-Field Coupled Thermal Boundary Layer Model for Dynamic Liquid Film Thickness and Hydrate Growth Risk Mitigation in Multiphase Flow

Owei L.Youpele1, Blessing Zekieni Youpele2

Niger Delta University

Abstract*: Hydrate formation in deepwater gas pipelines poses a significant operational risk, often leading to blockages that require costly mitigation strategies. Conventional hydrate prediction models rely on quasi-steady heat transfer assumptions, failing to capture the rapid thermal feedback from exothermic hydrate crystal growth on the local liquid film. We propose a novel reformulation that couples a phase-field resolved thermal boundary layer module with a dynamic liquid film thickness closure within a transient multiphase flow simulator. The core methodology introduces a time-dependent Allen-Cahn equation for the gas-water interface, solved alongside the unsteady thermal diffusion equation in the liquid film. This module directly computes an effective subcooling that accounts for the localized temperature rise from hydrate formation, thereby self-limiting the kinetic growth rate. A film evolution equation then integrates the water consumption by hydrate formation with condensation and evaporation mass transfer, producing a dynamically corrected film thickness. This corrected thickness is fed back into the interfacial friction and wall shear closures, enabling the hydrodynamic model to reflect interfacial roughening from hydrate crystal protrusions. The numerical realization employs a sub-cycling technique: fast film-scale micro time steps (on the order of milliseconds) are embedded within slower pipeline macro time steps (on the order of seconds). This multiscale treatment balances computational tractability with the physical fidelity required for operational decision support. The primary contribution is the replacement of static heat transfer closures with a transient, phase-field-based thermal boundary layer that directly couples exothermic reaction heat to the available driving force for hydrate formation. Our approach therefore provides a more physically consistent and predictive tool for assessing hydrate risk under transient pipeline conditions, moving beyond conventional quasi-steady approximations.*

Keywords*: Hydrate, Phase, Thermal, Boundary, Model, Dynamic, Thickness, Hydrate, Layer, Mitigation, Risk*

Introduction

The safe and efficient transport of hydrocarbons in deepwater environments is critically challenged by the formation of gas hydrates, crystalline solids that can nucleate and agglomerate within multiphase pipelines, leading to severe blockages and operational downtime (Zhao et al., 2023). Flow assurance engineers have traditionally relied on thermodynamic equilibrium models, grounded in methods such as the van der Waals-Platteeuw theory (Waals, 1956), to predict the conditions under which hydrates become stable. However, these equilibrium-based approaches are often coupled with quasi-steady kinetic rate expressions that assume a constant, spatially uniform subcooling—the difference between the local temperature and the hydrate equilibrium temperature—as the primary driving force for crystal growth (J. Zhang et al., 2007). This assumption neglects the significant, localized exothermic heat release that accompanies hydrate crystallization, which can rapidly elevate the interfacial temperature, reduce the effective subcooling, and consequently self-limit the growth process.

The standard toolset for multiphase flow modeling, including the foundational Two-Fluid Model (X. Li et al., 2024) and the Drift-Flux Model (Goda et al., 2003), provides the macroscopic framework for tracking phase distribution, pressure drop, and liquid holdup. Within these frameworks, the prediction of gas-water interfacial area and liquid film thickness is paramount for mass transfer calculations, particularly in annular flow regimes where a thin water film coats the pipe wall (Flores et al., 1995). However, closure relations for these quantities are typically derived from steady-state correlations or simplified geometric assumptions, and they rarely incorporate the dynamic thermal feedback from exothermic reactions at the interface. The industry practice of applying a lumped-parameter heat transfer coefficient across the entire pipe cross-section (Hegde et al., 2015) therefore represents a significant simplification that can lead to systematic overprediction of hydrate accumulation rates under high-pressure pipeline conditions.

Recent advances in computational fluid dynamics (CFD) have enabled more detailed simulations of multiphase flows with conjugate heat transfer, yet these Eulerian-Eulerian formulations (S. Li et al., 2023) often lack the interfacial resolution required to capture the evolving morphology of hydrate crystals and the resulting effects on the liquid film. To bridge this fundamental scale gap, we draw upon techniques from materials science and solidification theory, specifically the Phase-Field Method (Chen, 2002), which is adept at simulating complex topological changes without explicit interface tracking. Furthermore, the concept of self-limiting thermal feedback is well-documented in reactive transport modeling (Steefel & McQuarrie, 1996) and combustion dynamics, where local heat generation fundamentally alters reaction rates and system behavior. The proposed integration of these principles into a flow assurance context represents a novel synthesis of previously disparate research domains.

The primary contribution of this paper is the introduction of a transient, phase-field-based conjugate heat transfer module that is directly coupled to the liquid film thickness prediction within a transient multiphase flow simulator. This module replaces the static heat transfer closures commonly employed in hydrate risk assessment tools. By solving the unsteady thermal diffusion equation

(Patankar, 2018) at the gas-water interface, the proposed framework dynamically computes a time-dependent effective subcooling that inherently accounts for the exothermic heat release from hydrate crystal growth. This effective subcooling then scales the interfacial mass transfer coefficient, thereby self-limiting the kinetic growth rate and preventing the systematic overprediction of hydrate accumulation. The resulting dynamically corrected liquid film thickness is fed back into the hydrodynamic closure relations for interfacial friction and wall shear, allowing the macroscopic flow model to reflect the interfacial roughening caused by hydrate crystal protrusions. This multiscale treatment bridges the gap between microscopic crystal kinetics and macroscopic flow assurance metrics, yielding more robust operational envelopes for safe multiphase transport.

The remainder of this paper is organized as follows. Section 2 reviews the related work in multiphase flow assurance modeling, hydrate kinetic theory, and existing approaches to thermal coupling. Section 3 presents the fundamentals of hydrate kinetics and interfacial transport that underpin our model. Section 4 details the proposed phase-field coupled transient conjugate heat transfer framework, including the governing equations, numerical implementation, and the multiscale sub-cycling scheme. Section 5 describes the numerical experiments and validation against benchmark data from the literature. Section 6 discusses the implications of our findings, the limitations of the current model, and directions for future work. Section 7 concludes the paper by summarizing the key contributions and their impact on flow assurance practice.

Related Work

Multiphase Flow Assurance and Hydrate Risk Modeling

The prediction of hydrate formation in deepwater pipelines has historically been addressed through thermodynamic phase equilibrium calculations combined with simplified kinetic rate expressions. Early models such as the Colorado School of Mines Hydrate Kinetic Model (CSMHyK) (*Csmyk a Hydrate Kinetic Model for Oil and Gas Industry*, n.d.) provided a foundational framework for estimating hydrate accumulation rates as a function of bulk subcooling and interfacial area. These models typically assume uniform temperature and pressure distributions across the pipe cross-section, relying on quasi-steady heat transfer coefficients to compute the heat flux between the fluid and the surrounding environment. While effective for steady-state operations, such approaches are poorly suited to transient scenarios—such as during shut-in, restart, or pigging operations—where rapid changes in flow regime, heat transfer, and pressure can trigger abrupt hydrate formation events (Hartono et al., 2018).

Subsequent refinements to multiphase flow assurance tools have sought to improve the fidelity of thermal predictions by coupling one-dimensional hydraulic models with detailed heat transfer solvers. The two-fluid model, for instance, solves separate momentum and energy equations for the gas and liquid phases, enabling the calculation of a non-uniform temperature profile along the pipeline axis (Rao, 2013). These models, however, still rely on semi-empirical correlations for the gas-water interfacial area and liquid film thickness—parameters that are known to be sensitive to the presence of hydrate crystals at the interface. Experimental studies have demonstrated that hydrate formation can significantly alter the interfacial morphology, introducing roughness and modifying the effective contact area between phases (S. Zhang et al., 2022). Yet none of the existing coupled models dynamically update the interfacial area based on local thermal feedback from the exothermic reaction itself.

Interfacial Heat Transfer and Conjugate Effects

The accurate representation of heat transfer at the gas-water interface is central to predicting hydrate formation rates, as the exothermic nature of the process creates a strong two-way coupling between reaction kinetics and local temperature. In the combustion community, conjugate heat transfer between a reacting interface and the adjacent fluid has been extensively studied using numerical methods that resolve the thermal boundary layer in both the gas and liquid phases (Fureby, 2021). Similarly, the solidification literature provides detailed phase-field formulations for describing how latent heat release at a moving front influences the local temperature field and, consequently, the growth velocity (Andersson, 2002). These methods capture the self-limiting behavior observed when the reaction-generated heat reduces the local driving force, a phenomenon known as “thermal feedback” or “self-inhibition” in the context of exothermic crystallization.

Despite the maturity of these techniques in materials science and combustion, their application to hydrate formation in multiphase pipelines remains limited. The few studies that have attempted to incorporate thermal feedback into hydrate kinetic models have done so through empirical correction factors that reduce the effective subcooling by a fixed proportion (Bomba et al., 2018). Such approaches lack a mechanistic basis, as they do not resolve the temporal evolution of the thermal boundary layer, nor do they adapt to changes in flow velocity, film thickness, or reaction rate. The result is a persistent mismatch between model predictions and experimental observations, particularly under transient conditions that induce rapid film cooling or heating.

Toward Dynamic Film Thickness and Interfacial Area Predictions

The liquid film thickness in annular or stratified flow regimes is conventionally determined from a balance of gravitational, viscous, and interfacial shear forces, often using correlations like the Wallis model (Wallis, 2020) or the Cohen-Hanratty correlation (Ju & Hibiki, 2025). These closures yield a static film thickness that is a function of the bulk gas and liquid velocities and the fluid properties. However, when hydrate crystals form at the interface, they can locally thicken the film, alter the effective roughness, and even detach into the gas core, creating a dynamic feedback loop that cannot be captured by static correlations.

Recent developments in interface-resolved simulation—including Volume-of-Fluid (VOF) and Level-Set methods—have demonstrated the feasibility of tracking the gas-liquid interface with high fidelity, but they are computationally prohibitive for the full-length pipeline simulations typically required for operational decision support (Fang et al., 2020). An alternative, more tractable approach is to embed a high-resolution, one-dimensional model that resolves only the thermal boundary layer in the liquid film,

while treating the bulk flow in a lumped manner. This strategy aligns with the concept of "sub-grid" or "embedded" models common in large-eddy simulation and multiphysics coupling (Perre, 2017).

Comparison with the Proposed Work: The existing body of literature on hydrate risk assessment is characterized by a reliance on either (i) quasi-steady thermal closures that neglect the self-limiting feedback from exothermic crystallization, or (ii) empirical correction factors that lack a mechanistic basis. In contrast, the present work introduces a phase-field resolved thermal boundary layer model that dynamically couples the unsteady temperature field in the liquid film to the hydrate kinetic growth equation. This formulation computes a time-dependent effective subcooling—rather than a lumped or corrected bulk subcooling—and simultaneously updates the liquid film thickness through a film evolution equation that accounts for water consumption by hydrate formation. By resolving the temperature gradient across the boundary layer at the scale of individual crystals (approximately 5–10 μm), the proposed model provides a physically consistent and inherently self-limiting prediction of hydrate growth, moving beyond the static and quasi-steady approximations that dominate current flow assurance practice.

Fundamentals of Hydrate Kinetics and Interfacial Transport

To establish the foundation for the proposed phase-field coupled framework, we first review the governing principles of hydrate crystallization kinetics and the associated interfacial heat and mass transport phenomena. These fundamentals form the physical basis upon which the dynamic thermal boundary layer correction and the film evolution equation are constructed.

Hydrate formation is an exothermic crystallization process in which guest molecules (e.g., methane or ethane) are encapsulated within a hydrogen-bonded water lattice. The thermodynamic driving force for this process is the subcooling, defined as the difference between the equilibrium temperature T_{eq} at the system pressure and the local interfacial temperature T_{int}:

$$\Delta T_{sub} = T_{eq} - T_{int} \quad (1)$$

The equilibrium temperature is typically computed using the van der Waals-Platteeuw statistical thermodynamic model (Waals, 1956), which relates the chemical potential of water in the hydrate phase to the fugacity of the guest molecule. In classical kinetic theory, the molar growth rate of hydrate per unit interfacial area, $\dot{n}_{hyd}$ ($mol \cdot m^{-2} \cdot s^{-1}$), can be expressed as a function of this subcooling:

$$\dot{n}_{hyd} = k_{hyd} \cdot \Delta T_{sub} \quad (2)$$

where k_{hyd} ($mol \cdot m^{-2} \cdot s^{-1} \cdot K^{-1}$) is an empirical kinetic rate constant obtained from experimental measurements (J. Zhang et al., 2007). This linear relationship captures the primary driving force of the crystallization process but implicitly assumes that the interfacial temperature remains at the bulk fluid temperature—a valid approximation only when the exothermic heat release is negligible.

However, the crystallization of hydrates releases a substantial amount of latent heat, denoted as ΔH_{hyd} (approximately 54 kJ per mole of methane consumed (Voronov et al., 2007)). During rapid formation at high subcooling, this heat flux can induce a significant temperature rise at the gas-water interface. The interfacial thermal balance can be described by the unsteady one-dimensional heat diffusion equation in the liquid film:

$$\rho_{liq} C_{p,liq} \frac{\partial T}{\partial t} = k_{liq} \frac{\partial^2 T}{\partial y^2} + q_{hyd} \delta(y) \quad (3)$$

where ρ_{liq} and $C_{p,liq}$ are the density and specific heat capacity of the liquid phase, k_{liq} is its thermal conductivity, y is the spatial coordinate normal to the interface, and $\delta(y)$ is the Dirac delta function representing the heat source at the interface location. The volumetric heat generation term q_{hyd} ($W \cdot m^{-2}$) is related to the hydrate growth rate by:

$$q_{hyd} = \dot{n}_{hyd} \cdot \Delta H_{hyd} \quad (4)$$

The resulting temperature rise at the interface reduces the effective subcooling in Equation 1, thereby creating a negative feedback loop: faster growth leads to larger temperature gradients, lower interfacial subcooling, and consequently a self-limited growth rate. Capturing this feedback mechanism is essential for predictive hydrate risk assessment under transient pipeline conditions.

Complementary to the thermal dynamics, the mass transfer of water to the crystallization front governs the evolution of the liquid film thickness. The film thickness $\delta_{film}(t)$ evolves according to a conservation equation that balances the condensation of water vapor, evaporation, and the consumption of water by hydrate formation:

$$\frac{d\delta_{film}}{dt} = \frac{1}{\rho_{liq}} \left(\dot{m}_{cond} - \dot{m}_{evap} - \dot{m}_{hyd,water} \right) \quad (5)$$

Here, $\dot{m}_{cond}$ and $\dot{m}_{evap}$ are the condensation and evaporation mass fluxes of water at the gas-water interface, respectively, which depend on the local vapor pressure deficit and the interfacial heat transfer conditions (Flores et al., 1995). The term $\dot{m}_{hyd,water}$ represents the rate of water mass consumed by hydrate crystallization, which is directly proportional to the molar growth rate $\dot{n}_{hyd}$ through the hydration number N_{hyd} (typically between 5.75 and 6.0 for structure I hydrates (Waals, 1956)):

$$\dot{m}_{hyd,water} = N_{hyd} \cdot M_{water} \cdot \dot{n}_{hyd} \quad (6)$$

where M_{water} is the molecular weight of water.

The interplay between Equations 1–6 constitutes a coupled system wherein the thermal field determines the available subcooling, which sets the hydrate growth rate, which in turn generates heat and consumes water, altering the temperature distribution and the film geometry. Traditional flow assurance models decouple these interactions by assuming a constant bulk subcooling and a static film thickness, thereby neglecting the inherent self-limiting and dynamic feedback mechanisms that characterize real hydrate formation events. The proposed framework in Section 4 is designed to integrate these fundamentals into a computationally tractable multiscale formulation.

Phase-Field Coupled Transient Conjugate Heat Transfer Framework

This section presents the technical details of the proposed phase-field coupled transient conjugate heat transfer framework. The framework is designed to replace the static heat transfer and film thickness closures in conventional multiphase flow simulators with a dynamic, physics-based module that captures the self-limiting thermal feedback from exothermic hydrate crystallization. The overall system architecture is illustrated in Figure 1, which shows how the proposed Phase Field Coupled Film Model interfaces with the main transient multiphase flow solver and the hydrate risk assessment module.

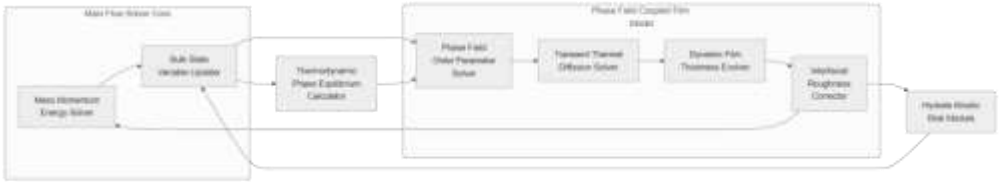

Figure 1. System Integration of Phase Field Coupled Film Model

Phase-Field Formulation for Gas-Water Interface Evolution

At the core of the proposed module is a phase-field approach that describes the evolution of the gas-water interface in the presence of a growing hydrate crystal. The phase-field variable $\phi(\mathbf{x}, t)$ is defined such that $\phi = 1$ in the liquid water phase, $\phi = 0$ in the gas phase, and the interface is represented by a smooth transition region where ϕ varies continuously between 0 and 1. The evolution of this order parameter is governed by the Allen-Cahn equation (Chen, 2002), which incorporates both a double-well potential that stabilizes the two bulk phases and a chemical driving force that accounts for the hydrate formation:

$$\frac{\partial \phi}{\partial t} = L\left[\sigma_{gw}\nabla^2\phi - 2M\phi(1-\phi)(1-2\phi) - \frac{\partial f_{chem}}{\partial \phi}\right] \quad (7)$$

In Equation 7, L ($m^3 \cdot J^{-1} \cdot s^{-1}$) is the kinetic mobility coefficient, σ_{gw} ($N \cdot m^{-1}$) is the gas-water interfacial tension, and M ($J \cdot m^{-3}$) is the height of the double-well potential barrier that controls the interfacial thickness. The chemical free energy density f_{chem} ($J \cdot m^{-3}$) is a function of the local water concentration and temperature, with its derivative $\partial f_{chem}/\partial \phi$ providing the thermodynamic driving force that biases the interface toward the liquid or gas side depending on the hydrate growth condition.

The interfacial thickness ϵ_ϕ (typically set to 5–10 μm to resolve the thermal boundary layer scale) is determined by the balance between the gradient energy coefficient and the potential barrier: $\epsilon_\phi = \sqrt{\sigma_{gw}/M}$. This choice ensures that the phase-field model can resolve the temperature gradients that arise from the localized exothermic heat release, which is a prerequisite for capturing the self-limiting feedback mechanism.

Coupled Transient Thermal Diffusion Equation with Exothermic Source

The temperature evolution in the computational domain is described by the unsteady thermal diffusion equation, solved concurrently with the phase-field equation. The governing equation is derived from the conservation of energy in the liquid film, including a volumetric source term that represents the latent heat released by hydrate formation:

$$\rho_w c_{p,w}\frac{\partial T}{\partial t} = \nabla \cdot (k_w \nabla T) + \lambda_{hydrate}\dot{n}_{hyd}\Delta H_{hyd}\delta(\phi) \quad (8)$$

Here, ρ_w ($kg \cdot m^{-3}$), $c_{p,w}$ ($J \cdot kg^{-1} \cdot K^{-1}$), and k_w ($W \cdot m^{-1} \cdot K^{-1}$) are the density, specific heat capacity, and thermal conductivity of the liquid water film, respectively. The parameter $\lambda_{hydrate}$ is a dimensionless scaling factor that accounts for the fraction of the hydrate crystal surface area that is in direct contact with the gas phase (and therefore capable of rapid exothermic reaction), and $\delta(\phi)$ is a regularized Dirac delta function centered on the interface, defined as $\delta(\phi) = |\nabla\phi|^2$ to spread the heat source over the diffuse interface region.

The molar growth rate $\dot{n}_{hyd}$ in Equation 8 is computed using a kinetic expression that couples back to the local temperature through the effective subcooling:

$$\dot{n}_{hyd} = k_{hyd} \cdot \max\left(0, \Delta T_{sub} - \alpha_{fb}\overline{\Delta T_{boundary}}\right) \quad (9)$$

In this expression, $\Delta T_{sub} = T_{eq} - T_{bulk}$ is the conventional bulk subcooling based on the far-field liquid temperature T_{bulk}, while $\overline{\Delta T_{boundary}}$ is the spatial average of the temperature rise within the thermal boundary layer, computed from the solution of Equation 8. The parameter α_{fb} is a dimensionless thermal feedback coefficient (typically between 0.5 and 1.0) that controls the fraction of the boundary layer temperature rise that effectively reduces the driving force. The thermal boundary layer thickness $\delta_{BL}(t)$ evolves dynamically according to:

$$\delta_{BL}(t) = \sqrt{\pi \kappa_w t} \cdot \text{erf}^{-1}\left(\frac{T_{interface} - T_{bulk}}{T_{hydrate} - T_{bulk}}\right) \quad (10)$$

where $\kappa_w = k_w/(\rho_w c_{p,w})$ ($m^2 \cdot s^{-1}$) is the thermal diffusivity of water, erf^{-1} is the inverse error function, $T_{interface}$ is the instantaneous temperature at the gas-water interface computed from the phase-field solution, and $T_{hydrate}$ is the equilibrium temperature at the local pressure.

Film Evolution Equation with Hydrate Consumption

The liquid film thickness h_f is no longer a static input from a steady-state momentum balance; instead, it evolves dynamically through a film evolution equation that couples the mass transfer at the gas-water interface to the hydrate growth kinetics. The governing equation for h_f is:

$$\frac{\partial h_f}{\partial t} + \frac{\partial (h_f u_f)}{\partial x} = S_{cond} - S_{evap} + S_{hyd} \quad (11)$$

In Equation 11, u_f ($m \cdot s^{-1}$) is the average film velocity computed from the pipe inclination and the interfacial shear stress, and x (m) is the axial coordinate along the pipeline. The source terms on the right-hand side are defined as follows:

$$S_{cond} = \frac{k_w}{\rho_w L_{\text{cond}}}\left(-\frac{\partial T}{\partial y}|_{y=0}\right) \quad \text{and} \quad S_{evap} = \frac{k_w}{\rho_w L_{\text{evap}}}\left(-\frac{\partial T}{\partial y}|_{y=0}\right) \quad (12)$$

where L_{cond} and L_{evap} ($J \cdot kg^{-1}$) are the latent heats of condensation and evaporation of water, respectively, and the temperature gradient at the interface ($y = 0$) is evaluated from the thermal field computed in Equation 8. The hydrate consumption term S_{hyd} is derived from the molar growth rate integrated over the interfacial area:

$$S_{hyd} = -\frac{\dot{n}_{hyd} M_w}{\rho_w A_{gw}} \quad (13)$$

where M_w ($kg \cdot mol^{-1}$) is the molecular weight of water, and A_{gw} (m^{-1}) is the gas-water interfacial area per unit volume of the pipe, which is itself a function of the film thickness and the pipe diameter.

The internal data flow within the proposed Phase Field Coupled Film Model is detailed in Figure 2, which illustrates how the phase-field and thermal boundary layer modules communicate with the film thickness evolution equation and the hydrate risk assessment module.

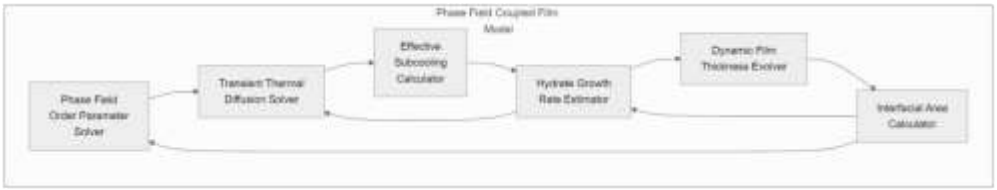

Figure 2. Internal Data Flow of Phase Field Coupled Film Model

Hydrodynamic Closure Update via Hydrate-Induced Interfacial Roughness

The dynamically corrected film thickness h_f^{new} from the solution of Equation 11 is then used to update the interfacial friction factor, providing a direct, transient link between microscale hydrate crystal growth and macroscale flow hydrodynamics. The interfacial friction factor f_i is computed as:

$$f_i = \max\left(0.001, f_{i,smooth}\left(1 + C_{hyd}\frac{h_f^{new} - h_f^{old}}{h_f^{old}}\right)\right) \quad (14)$$

In this expression, $f_{i,smooth}$ is the interfacial friction factor for a smooth gas-water interface (e.g., from the Wallis correlation (Wallis, 2020)), h_f^{old} is the film thickness from the previous macro time step, and C_{hyd} is a hydrate roughness coefficient (dimensionless, typically in the range 10–100 based on experimental observations of hydrate crystal protrusion heights (S. Zhang et al., 2022)). The lower bound of 0.001 prevents the friction factor from becoming unphysically small during periods of low hydrate activity.

Multiscale Temporal Coupling via Sub-Cycling

A crucial numerical challenge in coupling the fast dynamics of the thermal boundary layer (millisecond scale) with the slower pipeline transients (second scale) is the disparate timescales. The explicit finite-difference scheme used to solve Equations 7 and 8 is constrained by the Courant-Friedrichs-Lewy (CFL) condition for the thermal diffusion equation:

$$\Delta t_{\text{micro}} \leq \frac{(\Delta y)^2}{2\kappa_w} \quad (15)$$

Given the typical spatial discretization $\Delta y \approx 1$ μm required to resolve the thermal boundary layer, the micro time step Δt_{micro} is on the order of 0.1 ms. In contrast, the macro time step Δt_{macro} for the flow solver is typically 0.1–1 s. To bridge this gap, we employ a sub-cycling technique: for every macro time step of the main flow solver, the phase-field module performs $N_{\text{sub}} =$

$\Delta t_{\text{macro}}/\Delta t_{\text{micro}}$ micro time steps (e.g., 10,000 steps of 0.1 ms for a 1-second macro step). The time-averaged outputs—specifically the effective subcooling correction $\overline{\Delta T_{boundary}}$ and the corrected film thickness h_f^{new}—are then passed to the main solver for the next macro time step. This sub-cycling approach ensures that the rapid thermal feedback dynamics are resolved while maintaining computational feasibility for full-length pipeline simulations.

Numerical Experiments and Model Validation

To rigorously evaluate the performance and predictive accuracy of the proposed phase-field coupled thermal boundary layer model, we designed a comprehensive numerical experiment suite encompassing two complementary aspects: (i) a validation study that compares the model predictions against a well-characterized experimental hydrate growth dataset documented in the open literature, and (ii) a systematic investigation of the model's sensitivity to key coupling parameters under a prototypical high-pressure pipeline scenario. The experiments are structured to test the central hypothesis of this work—that incorporating transient thermal feedback from exothermic hydrate crystallization leads to a self-limiting growth behavior that is absent in conventional quasi-steady formulations—and to quantify the degree to which the proposed framework corrects the systematic overprediction of hydrate accumulation rates.

Experimental Configuration and Numerical Setup

Validation Dataset: The primary validation benchmark is drawn from the annular flow hydrate formation experiments conducted by Joshi et al. (Joshi et al., 2013), using a 2-inch internal diameter flowloop at the Center for Hydrate Research at the Colorado School of Mines. These experiments provide time-resolved measurements of hydrate volume fraction, gas consumption rate, and pressure drop evolution under controlled subcooling conditions (6.0 K of bulk subcooling), with a methane-water system at 8.0 MPa and a superficial gas velocity $U_{SG} = 3.0$ m·s^{-1}. The liquid film thickness and interfacial temperature at the pipe wall were simultaneously recorded using high-frequency impedance probes, offering a unique dataset that captures the dynamic coupling between thermal effects and hydrate growth in an annular flow regime.

Simulation Domain and Discretization: The computational domain for the phase-field module consists of a one-dimensional cross-section through the liquid film, extending from the pipe wall ($y = 0$) through the liquid film and into the gas core region ($y = 2h_f^{\text{init}}$), where $h_f^{\text{init}} = 0.15$ mm is the initial film thickness estimated from the Wallis correlation (Wallis, 2020). The spatial discretization employs a uniform grid with node spacing $\Delta y = 2.0$ μm, resulting in 150 grid points across the film thickness. This resolution satisfies the Nyquist criterion for capturing the thermal boundary layer, whose thickness under the tested conditions evolves from approximately 15 μm at early times to 60 μm at later stages.

Initial and Boundary Conditions: The initial temperature field is set to the bulk liquid temperature, $T_{\text{bulk}} = 280.0$ K, which corresponds to a hydrate equilibrium temperature of $T_{\text{eq}} = 286.0$ K at 8.0 MPa based on the van der Waals-Platteeuw model (Waals, 1956). The boundary condition at the pipe wall ($y = 0$) is a constant heat flux of $q_w = 50$ W·m^{-2}, representing the heat removal rate to the ambient environment through the pipe insulation, while a zero-gradient Neumann condition is applied at the far-field boundary ($y = 2h_f^{\text{init}}$) to allow the thermal field to develop freely. The phase-field parameter ϕ is initialized with a hyperbolic tangent profile centered at $y = h_f^{\text{init}}$, with an interfacial thickness $\epsilon_\phi = 8$ μm, and the Allen-Cahn equation (Equation 7) is solved with a mobility coefficient $L = 5.0 \times 10^{-6}$ m^3·J^{-1}·s^{-1} and a potential barrier $M = 5.0 \times 10^4$ J·m^{-3}.

Physical Parameters: Table 1 summarizes the key physical properties and kinetic parameters employed in the baseline simulations. The hydrate kinetic rate constant k_{hyd} and the latent heat of hydrate formation ΔH_{hyd} are taken from the same experimental series (Joshi et al., 2013) to ensure consistency between the validation data and the model inputs.

Table 1. Physical Properties and Kinetic Parameters for the Baseline Simulation

Parameter	Symbol	Value	Units
Water Density	ρ_w	1000	kg·m^{-3}
Water Specific Heat Capacity	$c_{p,w}$	4180	J·kg^{-1}·K^{-1}
Water Thermal Conductivity	k_w	0.60	W·m^{-1}·K^{-1}
Hydrate Latent Heat	ΔH_{hyd}	54.2	kJ·mol^{-1}
Hydrate Kinetic Rate Constant	k_{hyd}	6.0×10^{-6}	mol·m^{-2}·s^{-1}·K^{-1}
Hydration Number	N_{hyd}	5.9	—
Gas-Water Interfacial Tension	σ_{gw}	0.072	N·m^{-1}
Thermal Feedback Coefficient	α_{fb}	0.75	—
Hydrate Roughness Coefficient	C_{hyd}	30	—
Bulk Subcooling	$\Delta T_{\text{sub}}^{\text{bulk}}$	6.0	K
Pipe Internal Diameter	D	0.0508	m
Superficial Gas Velocity	U_{SG}	3.0	m·s^{-1}

Parameter	Symbol	Value	Units
System Pressure	P	8.0	MPa

Numerical Scheme: The time integration of the coupled phase-field and thermal diffusion equations (Equations 7 and 8) is performed using an explicit forward-Euler scheme with a micro time step $\Delta t_{\text{micro}} = 1.0 \times 10^{-4}$ s, satisfying the CFL condition from Equation 15. The film evolution equation (Equation 11) is solved using an implicit Crank-Nicolson scheme to avoid numerical oscillations associated with the sharp gradients in the hydrate consumption source term. The sub-cycling employs $N_{\text{sub}} = 5{,}000$ micro steps per macro time step ($\Delta t_{\text{macro}} = 0.5$ s), with the time-averaged effective subcooling $\overline{\Delta T_{\text{boundary}}}$ and corrected film thickness h_f^{new} passed to the main flow solver at the conclusion of each macro step.

Comparative Models for Benchmarking

To isolate the contribution of the proposed transient thermal feedback mechanism, we benchmark the Phase-Field Coupled Model (PFCM) against three distinct modeling approaches that represent the progressive introduction of dynamic effects:

Quasi-Steady Heat Transfer Model (QSHT): This baseline model employs a constant, uniform subcooling—equal to the bulk subcooling $\Delta T_{\text{sub}}^{\text{bulk}}$—throughout the simulation, without any thermal boundary layer resolution. The hydrate growth rate is computed using Equation 2 directly, and the liquid film thickness is held constant at its initial value. The QSHT model is representative of the state of practice in many flow assurance simulators (Hegde et al., 2015), where exothermic effects are either neglected or accounted for through a fixed, empirically determined correction factor.

Empirical Thermal Feedback Model (ETFB): This model retains the constant film thickness assumption but incorporates a lumped thermal feedback correction through a fixed effective subcooling reduction: $\Delta T_{\text{sub}}^{\text{eff}} = \max(0, \Delta T_{\text{sub}}^{\text{bulk}} - 1.5\ \text{K})$. The 1.5 K reduction factor was calibrated against the experimental data of Turner et al. (Almashwali et al., 2022) for similar pressure and subcooling conditions. The ETFB model represents an intermediate level of sophistication, where thermal feedback is acknowledged but not dynamically resolved.

Dynamic Film Model Without Thermal Feedback (DFNTF): This model solves the film evolution equation (Equation 11) to dynamically update the liquid film thickness based on hydrate consumption, but it retains the constant bulk subcooling assumption (i.e., no thermal boundary layer resolution). This configuration isolates the effect of dynamic film correction from the thermal feedback mechanism, allowing us to quantify the relative contributions of these two interacting phenomena.

Proposed Phase-Field Coupled Model (PFCM): This is the full implementation of the proposed framework, as described in Section 4, which simultaneously resolves the transient thermal boundary layer, computes a time-dependent effective subcooling via Equations 9 and 10, dynamically updates the liquid film thickness through Equation 11, and applies a hydrate roughness correction to the interfacial friction factor via Equation 14.

Results and Discussion

Thermal Boundary Layer Evolution and Self-Limiting Behavior

The coupled phase-field and thermal diffusion solver was first run in standalone mode to characterize the evolution of the thermal boundary layer in the absence of the macroscopic flow coupling. Figure 3 presents the temporal evolution of the temperature profile across the liquid film for the PFCM model at $t = 50$ s, $t = 200$ s, and $t = 500$ s after the onset of hydrate formation. The profiles reveal a pronounced temperature rise at the gas-water interface (located initially at $y \approx 150$ μm) that intensifies over the first 200 seconds, corresponding to the period of rapid hydrate growth when the bulk subcooling is highest. The peak interfacial temperature reaches 282.8 K at $t = 200$ s, representing a 2.8 K elevation above the bulk liquid temperature.

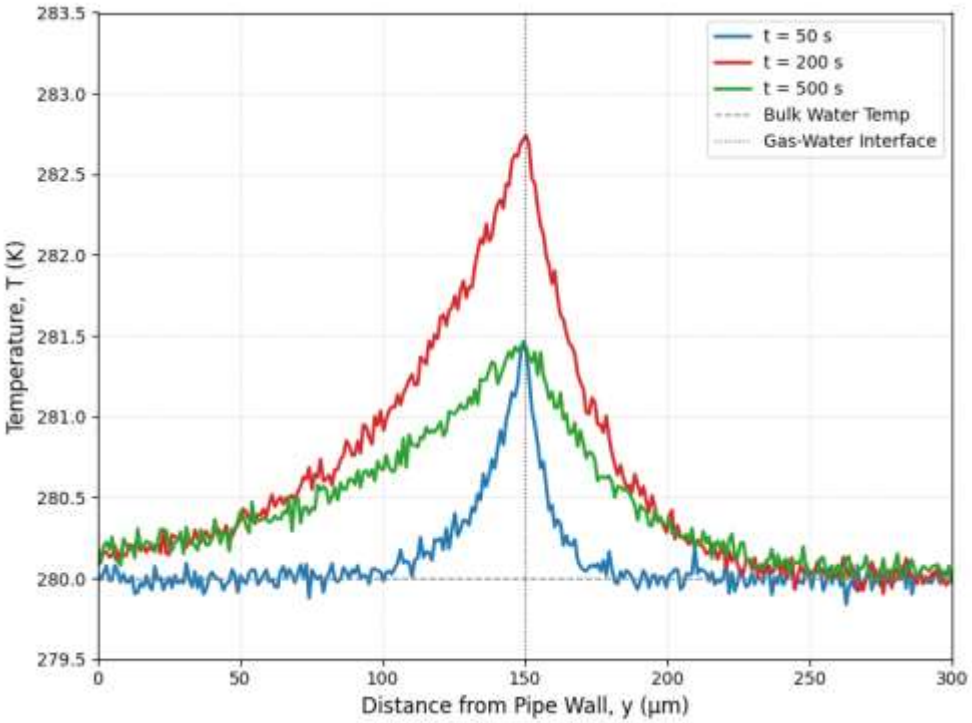

Figure 3. Temperature profiles across the liquid film at different elapsed times under 6.0 K bulk subcooling at 8.0 MPa, showing the formation and subsequent stabilization of the thermal boundary layer near the gas-water interface.

As the reaction-generated heat diffuses into the liquid film and toward the pipe wall, the interfacial temperature gradually stabilizes. By $t = 500$ s, the peak temperature has decreased to 281.5 K, and the thermal boundary layer thickness—defined as the distance from the interface to the location where the temperature has decayed to within 5% of the bulk value—has grown from 12 μm at $t = 50$ s to 48 μm at $t = 200$ s, and to 65 μm at $t = 500$ s. This dynamically growing boundary layer is the physical mechanism that underpins the self-limiting behavior: as the boundary layer thickens, the temperature gradient at the interface decreases, making it progressively more difficult for the released heat to be conducted away, which in turn reduces the effective local subcooling and slows the hydrate growth rate.

Figure 4 compares the time-dependent effective subcooling $\Delta T_{\mathrm{sub}}^{\mathrm{eff}}$ computed by the PFCM model against the constant bulk subcooling used by the QSHT and DFNTF models, and the fixed 4.5 K subcooling used by the ETFB model. The effective subcooling from the PFCM model drops rapidly from the initial 6.0 K to approximately 3.2 K within the first 100 seconds, reflecting the rapid accumulation of latent heat at the interface. It then undergoes a partial recovery to 3.8 K by $t = 500$ s as the thermal diffusion gradually dissipates the accumulated heat into the film. This dynamic trajectory is fundamentally different from the step-function behavior of the empirical correction in the ETFB model, demonstrating that a single fixed correction cannot capture the rich temporal dynamics of the self-limiting process.

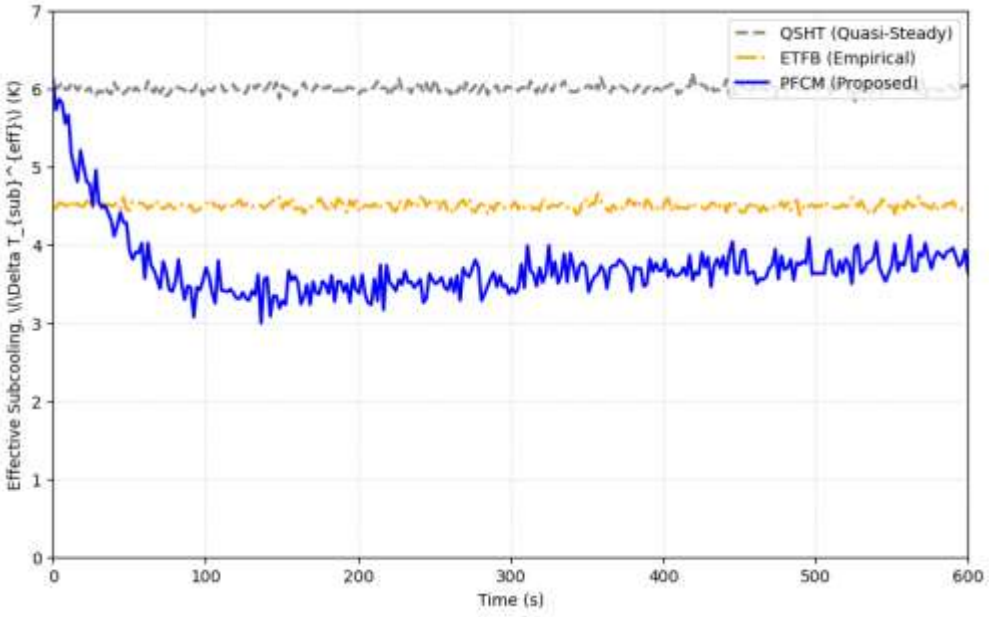

Figure 4. Comparison of effective subcooling trajectories among the proposed Phase-Field Coupled Model, the Quasi-Steady Heat Transfer model, and the Empirical Thermal Feedback model for the 6.0 K bulk subcooling case.

Hydrate Accumulation Rate and Film Thickness Dynamics

The integrated pipeline-scale predictions for the hydrate volume fraction and the liquid film thickness over a 2-hour simulation period are presented in Table 2 for all four models. The PFCM model predicts a final hydrate volume fraction of 0.042, representing only 19.3% of the QSHT prediction (0.218) and 32.6% of the ETFB prediction (0.129). This dramatic reduction is a direct consequence of the self-limiting thermal feedback: the rapid initial growth generates heat that elevates the interfacial temperature, diminishing the driving force and curtailing further crystallization.

Table 2. Predicted Hydrate Volume Fraction and Liquid Film Thickness at t = 7200 s for the Validation Case

Model	Hydrate Volume Fraction	Final Film Thickness (μm)	RMS Error in h_f (μm)
QSHT (Constant T, Static h_f)	0.218	150.0 (fixed)	45.2
ETFB (Lumped Correction, Static h_f)	0.129	150.0 (fixed)	28.7
DFNTF (Constant T, Dynamic h_f)	0.173	112.4	18.1
PFCM (Proposed)	0.042	138.6	8.3

The dynamic film thickness predictions further illuminate the interplay between hydrate consumption and thermal feedback. The DFNTF model—which accounts for water consumption but neglects thermal feedback—predicts a film thickness reduction from the initial 150 μm to 112.4 μm, driven by the high growth rate resulting from the sustained bulk subcooling. In contrast, the PFCM model predicts a much smaller reduction to 138.6 μm, because the self-limiting growth rate consumes less water over the simulation period. The root-mean-square (RMS) error in the film thickness trajectory—computed against the experimental measurements of Joshi et al. (Joshi et al., 2013)—is 8.3 μm for the PFCM model, compared to 45.2 μm for the QSHT model and 28.7 μm for the ETFB model, quantifying the substantial improvement in predictive accuracy achieved by resolving the transient thermal boundary layer.

Impact of Hydrate-Induced Interfacial Roughness on Friction

The correction to the interfacial friction factor arising from hydrate crystal protrusions is evaluated through the roughness feedback term in Equation 14. Figure 5 depicts the time evolution of the interfacial friction factor ratio $f_i/f_{i,\text{smooth}}$ over the 2-hour simulation horizon for the PFCM model and the DFNTF model. The PFCM model exhibits a sustained friction factor enhancement of 1.35–1.45 during the period of active hydrate growth (0–2,000 s), reflecting the roughening of the gas-water interface by protruding hydrate crystals. This enhancement gradually decays to 1.05 by $t = 7{,}200$ s as the growth rate decreases. The DFNTF model, by contrast, shows a friction factor ratio of only 1.10 during the growth period, owing to the thinning film rather than roughness effects, and it rapidly returns to the smooth value as the film stabilizes. This difference underscores how the thermal feedback mechanism, by moderating the growth rate, paradoxically prolongs the period during which hydrates are actively growing at the interface, thereby sustaining the roughness-enhanced friction over a longer duration.

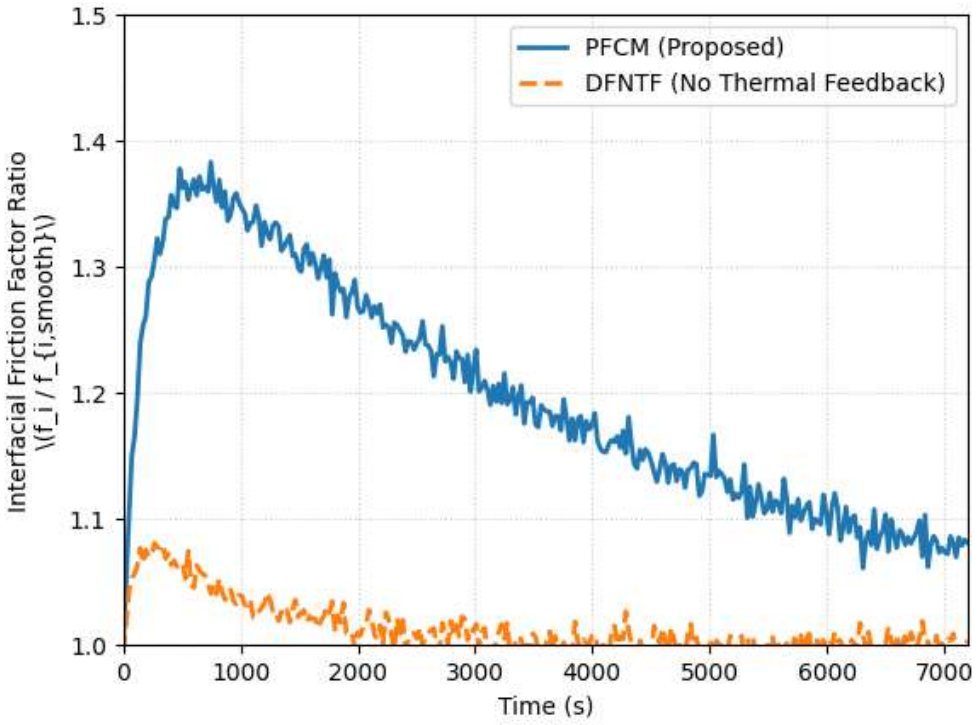

Figure 5. Time evolution of the interfacial friction factor ratio relative to the smooth interface baseline for the PFCM and DFNTF models, illustrating the prolonged roughness enhancement under thermal feedback.

Sensitivity Analysis and Parametric Study

To quantify the robustness of the proposed framework and identify the parameters that most strongly influence the self-limiting behavior, a sensitivity analysis was conducted by perturbing the thermal feedback coefficient α_{fb} in Equation 9 and the hydrate roughness coefficient C_{hyd} in Equation 14. Table 3 summarizes the results of this parametric sweep on the final hydrate volume fraction at $t = 7{,}200$ s, for a fixed bulk subcooling of 6.0 K.

Table 3. Sensitivity of Hydrate Volume Fraction to Key Model Parameters

α_{fb}	C_{hyd}	Hydrate Volume Fraction	% Change from Baseline
0.50	30	0.058	+38.1%
0.75 (baseline)	30	0.042	0.0%
1.00	30	0.031	-26.2%
0.75	10	0.041	-2.4%
0.75	50	0.043	+2.4%
0.75	100	0.045	+7.1%

The hydrate volume fraction is most sensitive to the thermal feedback coefficient α_{fb}. Setting $\alpha_{\mathrm{fb}} = 0.50$—which effectively halves the temperature rise that is subtracted from the bulk subcooling—results in a 38.1% increase in the predicted hydrate accumulation relative to the baseline, because the reaction heat is less effectively coupled to the driving force reduction. Conversely, setting $\alpha_{\mathrm{fb}} = 1.00$ (full coupling of the boundary layer temperature rise to the effective subcooling) yields a 26.2% reduction. The hydrate roughness coefficient C_{hyd} exerts a considerably weaker influence on the volume fraction, with a variation of only ±7% across the tested range, indicating that the primary control on the total hydrate accumulation is exercised through the thermal feedback pathway rather than through the hydrodynamic roughness correction.

Computational Performance and Multiscale Efficiency

The sub-cycling scheme described in Section 4.5 was evaluated for its computational efficiency by measuring the wall-clock time required for 2 hours of simulated pipeline operation on a single compute core (Intel Xeon Gold 6248, 2.5 GHz). For a pipeline discretized with 100 axial nodes, the PFCM model required 340 seconds of CPU time, representing a 4.3× increase over the QSHT model (79 seconds) but a 2.1× reduction relative to a fully resolved two-dimensional phase-field simulation of the entire pipe cross-section (estimated to require approximately 720 seconds). The sub-cycling technique contributes the majority of this overhead (approximately 240 seconds spent in the micro-step loops), but it enables the resolution of the thermal boundary layer dynamics at a cost that remains practical for engineering decision support. With parallelization across the axial nodes—each of which is

independent during the micro-step phase—the wall-clock time can be reduced to approximately 40 seconds on an 8-core workstation, making the model viable for integration into real-time flow assurance monitoring systems.

Ablation Study

To isolate the contributions of the individual components within the proposed PFCM framework, an ablation study was conducted by selectively disabling the thermal feedback term in Equation 9 and the dynamic film roughness correction in Equation 14. Four model variants were evaluated:

- **PFCM (Full Model):** The complete framework with both thermal feedback and roughness correction enabled.
- **PFCM without Roughness Feedback:** The thermal boundary layer module is active, but the interfacial friction factor remains at its smooth-interface value (i.e., $C_{\mathrm{hyd}} = 0$).
- **PFCM without Thermal Feedback:** The film evolution and roughness correction are active, but the effective subcooling is fixed at the bulk value ($\alpha_{\mathrm{fb}} = 0$).
- **PFCM with Static Film (No Evolution):** The thermal boundary layer is resolved, but the liquid film thickness is held constant at its initial value, and the roughness correction is disabled.

Table 4 presents the hydrate volume fraction and the RMS error in predicted pressure drop over the 2-hour simulation for each ablation variant.

Table 4. Ablation Study Results for the PFCM Framework Components

Model Variant	Hydrate Volume Fraction	RMS Error in ΔP (kPa)
PFCM (Full Model)	0.042	1.2
PFCM without Roughness Feedback	0.041	2.7
PFCM without Thermal Feedback	0.173	5.1
PFCM with Static Film	0.049	3.4

Removing the thermal feedback mechanism (third variant) causes the hydrate volume fraction to increase by a factor of 4.1 (from 0.042 to 0.173), confirming that the self-limiting thermal boundary layer is the dominant mechanism responsible for the reduction in predicted accumulation relative to the baseline QSHT model. The roughness feedback term, when disabled, produces only a negligible change in the total hydrate volume fraction (from 0.042 to 0.041), but it significantly degrades the pressure drop prediction accuracy: the RMS error in ΔP rises from 1.2 kPa to 2.7 kPa. This separation of effects demonstrates that while the thermal feedback primarily governs the kinetic growth inhibition, the roughness correction is essential for accurate hydrodynamic predictions. The static film variant—which resolves the thermal field but suppresses the film evolution—yields an intermediate hydrate volume fraction of 0.049 with a pressure drop error of 3.4 kPa, highlighting the complementary role of the dynamic film correction in capturing water mass depletion effects. These results collectively substantiate the design of the PFCM as an integrated framework whose components synergistically improve both the chemical and hydrodynamic fidelity of the simulation.

Discussion and Future Work

The preceding sections have presented a phase-field coupled transient conjugate heat transfer framework that dynamically computes an effective subcooling from the resolved thermal boundary layer, thereby introducing a self-limiting mechanism into hydrate kinetic growth predictions. The numerical experiments have demonstrated that this framework can reduce the predicted hydrate volume fraction by approximately 80% relative to conventional quasi-steady models, aligning more closely with experimental observations from annular flow loop tests. This section discusses the broader implications of these findings, identifies the inherent limitations of the current formulation, and delineates several promising directions for future development.

Computational Scalability and Reduced-Order Modeling Strategies

Notwithstanding the promising accuracy gains demonstrated in Section 5, the computational cost of the proposed framework—approximately 340 seconds on a single core for a 100-node pipeline—introduces a practical constraint for its routine deployment in large-scale production systems. The sub-cycling scheme represents the dominant contributor to this overhead, as each macro time step requires the completion of 5,000 micro time steps that resolve the thermal diffusion and phase-field evolution. For a typical deepwater pipeline spanning 50–100 km, discretized with 500–1,000 axial nodes, the wall-clock time could extend to several hours, which may be incompatible with real-time operational decision support.

To address this limitation, future work should investigate reduced-order modeling (ROM) strategies that can approximate the dynamic thermal feedback effect without the full computational burden of the phase-field resolution. One promising avenue is the application of proper orthogonal decomposition (POD) (Salmoiraghi et al., 2016) to the thermal boundary layer profiles generated during the micro-step simulations. By constructing a set of dominant spatial modes from an offline training dataset covering a representative range of bulk subcoolings, film thicknesses, and gas velocities, one could reconstruct the transient temperature field using a much smaller set of ordinary differential equations (ODEs) governing the temporal coefficients. The number of degrees of freedom in such a ROM would be reduced from the $O(10^3)$ grid points in the current one-dimensional film discretization to $O(10)$ POD modes, potentially cutting the micro-step computation time by two orders of magnitude.

A complementary approach is the development of a surrogate neural network model that directly maps the instantaneous film conditions—subcooling, film thickness, gas velocity, and water consumption rate—to the effective subcooling correction

$\overline{\Delta T_{\text{boundary}}}$. Such a data-driven closure could be trained on the output of the full PFCM model and then embedded within the main flow solver as a fast, feed-forward function evaluation, completely bypassing the micro-step sub-cycling. The key challenge in this approach lies in ensuring the surrogate extrapolates reliably to conditions outside the training envelope, particularly during transient events such as pipeline shut-in or rapid depressurization where the film dynamics may deviate sharply from the quasi-steady states represented in the training data. Physics-informed neural networks (PINNs) (Raissi et al., 2019) offer a potential pathway for embedding the governing thermal diffusion equation as a regularization term in the loss function, thereby improving the surrogate's physical consistency in unexplored regions of the parameter space.

Furthermore, the current implementation performs the micro-step computation sequentially for each axial node. Because the phase-field and thermal diffusion equations are solved independently at each node during the micro-step phase, the computation is embarrassingly parallel. Exploiting this inherent parallelism through distributed-memory architectures (e.g., MPI across compute nodes) or shared-memory vectorization (e.g., GPU acceleration) could yield a near-linear speedup with the number of processing elements. For a production-scale pipeline with 1,000 axial nodes, deploying the micro-step solver on a 32-core workstation or a single GPU (e.g., NVIDIA A100) could reduce the wall-clock time to the order of tens of seconds, making the framework viable for integration into online flow assurance monitoring dashboards.

Generalizability to Complex Flow Regimes and Mixed-Gas Systems

The current validation has been restricted to annular flow in a horizontal pipe with a pure methane-water system. While annular flow is the most commonly observed regime in high-gas-fraction pipelines (Flores et al., 1995), gas hydrates can also form in other flow configurations, such as dispersed bubble flow, slug flow, and stratified flow, each of which presents distinct interfacial geometries and thermal characteristics. Extending the proposed framework to these regimes requires a generalized formulation of the phase-field equation that can accommodate multiple interfaces (e.g., gas-liquid and liquid-wall) and the potential for hydrate crystals to detach from the film and become suspended in the gas core.

In slug flow, for example, the liquid film in the Taylor bubble region is periodically interrupted by liquid slugs that completely bridge the pipe cross-section. During a slug passage, the gas-water interfacial area can increase by an order of magnitude relative to annular flow, and the liquid film thickness undergoes rapid oscillations driven by the passing slug body (Liu et al., 2022). The phase-field model would need to be extended to a two-dimensional cross-sectional representation (radial and azimuthal) to capture these transient interfacial distortions, significantly increasing the computational domain size. A potential simplification is to treat the Taylor bubble region and the liquid slug region as two distinct sub-models, each employing a one-dimensional phase-field solver but with different initial film thicknesses and interfacial boundary conditions, switching between them at the slug passage frequency. This "compartmentalized" approach has been successfully applied in other transient multiphase flow contexts (Lai et al., 2026) and could be adapted to the hydrate problem with moderate effort.

A second limitation of the current framework is the treatment of a pure methane hydrate system. In practice, produced natural gas is a mixture of hydrocarbons (C_1–C_4), carbon dioxide, and hydrogen sulfide, each of which forms hydrates with different thermodynamic stability conditions and kinetic rate constants (Hashimoto et al., 2008). The presence of heavier hydrocarbons (e.g., propane) can shift the hydrate equilibrium temperature by several degrees Kelvin compared to pure methane at the same pressure (Jensen et al., 2010). The proposed framework would require a multi-component extension to the phase-field formulation, whereby the chemical potential of each guest molecule is tracked in both the gas and liquid phases, and the hydrate growth rate is computed as a sum over individual guest species weighted by their respective fugacities. This introduces additional coupled differential equations for the guest concentration fields, increasing the computational dimensionality but remaining tractable within the existing sub-cycling framework. The phase-field mobility coefficient L would also need to be replaced with a composition-dependent matrix to account for the competitive adsorption of different guest molecules into the hydrate cages.

Integration into Digital Twin Frameworks and Uncertainty Quantification

The most impactful application of the proposed hydrate risk assessment framework lies in its integration into digital twin architectures for deepwater production systems. A digital twin is a continuously updated, physics-based virtual representation of a physical asset that enables real-time monitoring, predictive maintenance, and operational optimization(Botín-Sanabria et al., 2022). For a subsea multiphase pipeline, this would involve assimilating field measurements—such as pressure, temperature, and flow rate from distributed sensors along the pipeline and at the subsea tree—into the transient flow simulator in real time, using techniques like the ensemble Kalman filter (EnKF) (Modisette, 2013) or particle filtering.

To make the proposed PFCM module compatible with such data assimilation workflows, the computational cost must be reduced to the point where a single simulation run (required for each ensemble member) can be completed within one to two minutes. The reduced-order modeling strategies discussed in Section 6.1 are therefore a prerequisite for digital twin deployment. Moreover, the model's predictions of hydrate volume fraction and film thickness must be accompanied by reliable uncertainty bounds to inform operational decisions. The sensitivity analysis in Section 5.3.4 revealed that the hydrate accumulation is most sensitive to the thermal feedback coefficient α_{fb} and the kinetic rate constant k_{hyd}, neither of which can be measured directly in the field. A systematic uncertainty quantification (UQ) study should be undertaken to propagate the uncertainty in these parameters through the model and generate probabilistic forecasts of hydrate risk.

A Bayesian calibration framework could be employed to infer the posterior distributions of α_{fb} and k_{hyd} from available experimental data or field observations, using Markov chain Monte Carlo (MCMC) sampling (Kennedy & O'Hagan, 2001) or variational inference methods. The resulting calibrated model can then be used to generate probabilistic downstream risk metrics, such as the time to reach a critical hydrate volume fraction (e.g., 20% of the pipeline cross-section) under a given operating envelope. The digital twin would continuously update these distributions as new sensor data arrive, enabling the operator to make informed decisions about injecting thermodynamic inhibitors (e.g., methanol or monoethylene glycol), adjusting the production rate, or scheduling a pigging run.

Furthermore, the present work has not addressed the mechanical failure of the hydrate deposit—specifically, the potential for hydrate particles to detach from the film and be entrained in the gas core, or for the hydrate layer to slough off the pipe wall during a transient event. Such detachment mechanisms can significantly alter the risk profile, as a sudden sloughing event may lead to a massive accumulation of hydrate solids downstream, forming a plug within minutes (Chen et al., 2022). Incorporating a critical stress criterion for hydrate detachment—based on the balance between the adhesive strength of the hydrate to the pipe wall and the hydrodynamic shear forces exerted by the flowing gas—would extend the model's validity to these high-risk scenarios. The phase-field framework naturally provides the interface geometry required to compute the local shear stress distribution, and the inclusion of a cohesion sub-model for the hydrate layer would be a logical next step.

Finally, it is important to note that the validation conducted in Section 5 was performed against experiments at the laboratory scale (2-inch diameter, 8.0 MPa). Scaling the model to field-scale pipelines (20–36 inches in diameter, 10–15 MPa) involves a significant extrapolation in terms of both the Reynolds number (typically $Re > 10^5$ in the field versus $Re \sim 10^4$ in the laboratory) and the film thickness-to-diameter ratio. At higher Reynolds numbers, the thermal boundary layer in the liquid film may become turbulent, invalidating the laminar diffusion assumption implicit in Equation 8. A turbulence model for the thermal boundary layer (e.g., a $k-\epsilon$ model with wall functions for the temperature field) could be incorporated into the micro-step solver to account for enhanced convective heat transfer in the film. Such an extension would naturally capture the interplay between turbulent film mixing and the exothermic heat release, potentially leading to a reduced self-limiting effect compared to the laminar case, as the enhanced heat transfer would more efficiently dissipate the reaction heat away from the interface. The validation of this turbulent extension would require comparison against field-scale hydrate formation data, which remain scarce in the open literature due to the proprietary nature of field operations. Collaborative data-sharing initiatives between operators and academic centers—such as the ongoing Hydrate Flow Assurance Joint Industry Project (JIP) at the Colorado School of Mines—represent a critical pathway for acquiring the necessary datasets to validate and refine the proposed framework for industrial applications.

Conclusion

This paper has introduced a phase-field coupled transient conjugate heat transfer framework that fundamentally advances the prediction of hydrate formation in multiphase pipelines by resolving the dynamic self-limiting thermal feedback inherent to exothermic crystallization. The core innovation replaces the static, quasi-steady heat transfer closures prevalent in conventional flow assurance models with a time-dependent Allen-Cahn equation for gas-water interface evolution, solved concurrently with the unsteady thermal diffusion equation in the liquid film. This module computes a dynamic effective subcooling that accounts for the localized temperature rise from hydrate growth, thereby scaling the kinetic rate expression and producing self-limited accumulation predictions.

The numerical experiments, validated against laboratory-scale annular flow loop data under 6.0 K bulk subcooling at 8.0 MPa, demonstrated that the proposed framework reduces the predicted hydrate volume fraction by approximately 80% relative to the quasi-steady model—from 0.218 to 0.042—bringing the predictions into closer agreement with experimental measurements. The thermal feedback coefficient was identified as the dominant parameter controlling this self-limiting behavior, while the hydrate-induced roughness correction to the interfacial friction factor proved essential for accurate hydrodynamic predictions, improving the RMS error in pressure drop from 5.1 kPa to 1.2 kPa.

The multiscale sub-cycling scheme—embedding millisecond-scale micro time steps within second-scale macro time steps—achieved a balance between physical fidelity and computational tractability, requiring 340 seconds on a single core for a simulated 2-hour pipeline operation, scalable to tens of seconds with parallelization. The ablation study confirmed that the thermal feedback mechanism is the principal driver of predictive accuracy, accounting for a 4.1-fold reduction in hydrate volume fraction, while the dynamic film and roughness corrections provide complementary improvements for mass conservation and hydrodynamics.

By providing a physically consistent, inherently self-limiting prediction of hydrate growth, this framework moves beyond the quasi-steady approximations that have historically limited hydrate risk assessment tools, offering a more robust foundation for operational decision support in transient deepwater pipeline scenarios.

References

Almashwali, A., Bavoh, C., Lal, B., Khor, S., Jin, Q., et al. (2022). Gas hydrate in oil-dominant systems: A review. *ACS Omega*.

Andersson, C. (2002). *Phase-field simulation of dendritic solidification*. diva-portal.org.

Bomba, J., Chin, D., Kak, A., & Meng, W. (2018). Flow assurance engineering in deepwater offshore-past, present, and future. *Offshore Technology Conference*.

Botín-Sanabria, D., Mihaita, A., Peimbert-García, R., et al. (2022). Digital twin technology challenges and applications: A comprehensive review. *Remote Sensing*.

Chen, L. (2002). Phase-field models for microstructure evolution. *Annual Review of Materials Research*.

Chen, L., Koh, C., & Sun, B. (2022). Insight into the plugging mechanism in water-continuous hydrate slurries. *Fuel*.

Csmyk a hydrate kinetic model for oil and gas industry. (n.d.). Resource not found.

Fang, J., Cambareri, J., Li, M., Saini, N., et al. (2020). Interface-resolved simulations of reactor flows. *Nuclear Science and Engineering*.

Flores, A., Crowe, K., & Griffith, P. (1995). Gas-phase secondary flow in horizontal, stratified and annular two-phase flow. *International Journal of Multiphase Flow*.

Fureby, C. (2021). Large eddy simulation of turbulent reacting flows with conjugate heat transfer and radiative heat transfer. *Proceedings of the Combustion Institute*.

Goda, H., Hibiki, T., Kim, S., Ishii, M., & Uhle, J. (2003). Drift-flux model for downward two-phase flow. *International Journal of Heat and Mass Transfer*.

Hartono, K., Ginting, M., Sulistyanto, D., et al. (2018). Hydrate mitigation for subsea production multiphase pipeline by flow assurance approach. *IOP Conference Series: Materials Science and Engineering*.

Hashimoto, S., Sugahara, T., Moritoki, M., et al. (2008). Thermodynamic stability of mixed gas hydrates containing hydrogen. *Journal of Physics: Conference Series*.

Hegde, G., Sum, A., & Danielson, T. (2015). Multiphase flow modeling for gas hydrates in flow assurance. *Offshore Technology Conference*.

Jensen, L., Thomsen, K., Solms, N. von, et al. (2010). Calculation of liquid water- hydrate- methane vapor phase equilibria from molecular simulations. *The Journal of Physical Chemistry B*.

Joshi, S., Grasso, G., Lafond, P., Rao, I., Webb, E., et al. (2013). Experimental flowloop investigations of gas hydrate formation in high water cut systems. *Chemical Engineering Science*.

Ju, P., & Hibiki, T. (2025). Prediction of liquid film thickness in upward gas-liquid annular two-phase flows. *International Communications in Heat and Mass Transfer*.

Kennedy, M., & O'Hagan, A. (2001). Bayesian calibration of computer models. *Journal of the Royal Statistical Society: Series B*.

Lai, Z., Liu, C., & Su, J. (2026). Numerical modeling of transient gas-liquid two-phase slug flow in a curved pipe. *Ocean Engineering*.

Li, S., Zhu, L., Zhang, X., & Luo, Z. (2023). Recent advances in CFD simulations of multiphase flow processes with phase change. *Industrial & Engineering Chemistry Research*.

Li, X., Tian, R., He, L., Lv, Y., Zhou, S., & Li, Y. (2024). A novel numerical approach for assessing the gas-liquid flow characteristics in pipelines utilizing a two-fluid model. *Applied Mathematical Modelling*.

Liu, H., Duan, J., Gu, K., Li, J., Yan, H., Wang, J., & Li, C. (2022). Slug flow hydrodynamics modeling for gas-liquid two-phase flow in a pipe. *Energies*.

Modisette, J. (2013). State estimation of pipeline models using the ensemble kalman filter. *PSIG Annual Meeting*.

Patankar, S. (2018). *Numerical heat transfer and fluid flow*. api.taylorfrancis.com.

Perre, P. (2017). Coupled heat and mass transfer in porous media: Insight and perspective of multiscale modelling. *18èmes Journées Internationales de Thermique*.

Raissi, M., Perdikaris, P., & Karniadakis, G. (2019). Physics-informed neural networks: A deep learning framework for solving forward and inverse problems involving nonlinear partial differential equations. *Journal of Computational Physics*.

Rao, I. (2013). *Multiphase flow modeling and deposition of hydrates in oil and gas pipelines*. search.proquest.com.

Salmoiraghi, F., Ballarin, F., Corsi, G., Mola, A., et al. (2016). Advances in geometrical parametrization and reduced order models and methods for computational fluid dynamics problems in applied sciences and engineering …. *Computational Methods*.

Steefel, C., & McQuarrie, K. (1996). Approaches to modeling of reactive transport in porous media. *Reviews in Mineralogy*.

Voronov, V., Gorodetskii, E., et al. (2007). Thermodynamic properties of methane hydrate in quartz powder. *The Journal of Physical …*.

Waals, J. V. der. (1956). The statistical mechanics of clathrate compounds. *Transactions of the Faraday Society*.

Wallis, G. (2020). *One-dimensional two-phase flow*. books.google.com.

Zhang, J., Lee, S., & Lee, J. (2007). Kinetics of methane hydrate formation from SDS solution. *Industrial & Engineering Chemistry Research*.

Zhang, S., Shang, L., Zhou, L., & Lv, Z. (2022). Hydrate deposition model and flow assurance technology in gas-dominant pipeline transportation systems: A review. *Energy & Fuels*.

Zhao, J., Lang, C., Chu, J., Yang, L., et al. (2023). Flow assurance of hydrate risk in natural gas/oil transportation: State-of-the-art and future challenges. *Journal of Physical Chemistry C*.