

Original Research

Real-Time Probabilistic Inference of Influx Dynamics Using Coupled Thermo-Solubility and Two-Phase Flow Models

Nguyen Hoang Minh¹ and Tran Quoc Dat²

¹Dong Nai Technology University, 206 Nguyen Khuyen Road, Bien Hoa City, Vietnam.

²An Giang University, 18 Ung Van Khiem Street, Long Xuyen City, Vietnam.

Abstract

Gas influx events during drilling remain a persistent source of operational risk because the observable signals used for detection and response are indirect, delayed, and confounded by uncertain multiphase physics. Surface pressure and flow measurements provide high-availability telemetry, yet they are filtered through annular hydraulics, compressibility, and regime-dependent slip, while the downhole state of the influx evolves through dissolution, expansion, and migration. This paper develops a real-time probabilistic digital-twin framework for kick detection and response that explicitly couples thermo-solubility, transient two-phase transport, and regime-switching closures with sequential Bayesian inference and risk-aware control. The central technical contribution is a hybrid reduced-order model whose continuous state captures pressure, holdup, dissolved-gas inventory, and thermal effects, while a discrete latent variable encodes flow-regime-dependent slip and friction structure. The model is embedded in an online estimator that assimilates surface pressure and flow data to infer posterior distributions over influx size, location, and composition-sensitive transport parameters, with quantified uncertainty and detectability metrics. A risk-sensitive decision layer then maps the inferred state to choke and circulation policies formulated as chance-constrained optimal control, with constraints representing pressure integrity margins and operational actuation limits. The resulting architecture targets field utility by enforcing computational budgets, bounded-error updates, and stability of closed-loop decisions under noisy telemetry. The framework is intended as a principled bridge between mechanistic hydraulics and data-driven signal interpretation, emphasizing deployable inference and control rather than retrospective reconstruction.

1. Introduction

Managing unintended formation inflow during drilling is fundamentally an inference-and-control problem under partial observability. The wellbore is a long, compliant, thermally stratified conduit in which small changes in downhole state can be attenuated, delayed, or transformed before they appear at the surface. In practice, the decision maker must map a stream of surface signals into an actionable hypothesis about whether an influx exists, where it is, how large it is, and which intervention is safest given competing constraints such as avoiding further underbalance, preventing excessive surface pressures, and respecting fracture limits. The difficulty is not solely the complexity of multiphase flow, but also the fact that the information content of surface measurements depends on transient dynamics, sensor placement, and the regime-dependent relationship between pressure gradient and holdup. A modern well-control tool must therefore treat kick handling as a closed-loop system: a physics model that is fast enough for online use, an estimator that quantifies uncertainty rather than returning a single deterministic reconstruction, and a decision layer that accounts for safety margins probabilistically rather than through ad hoc conservatism.

A second challenge is model mismatch. High-fidelity numerical simulators can represent detailed annular transport and thermodynamics, but their computational cost and parameter demands often prevent deployment under tight operational latencies [1]. Conversely, simplified hydraulics formulas can be

run quickly but typically rely on closures whose validity can shift abruptly with changing flow regime, mud rheology, temperature, and gas dissolution behavior. The consequence is a gap between tractable real-time computation and physically credible prediction. Bridging this gap calls for reduced-order modeling that preserves the mechanisms that matter for observability and control, while abstracting away degrees of freedom that cannot be identified from available telemetry. The critical point is that reduction must be done with the inference problem in mind: a reduced model is useful only insofar as it yields a stable mapping from surface signals to latent influx state under realistic noise and parameter uncertainty.




Figure 1: Online digital-twin loop for gas-influx handling: surface measurements are synchronized and filtered, assimilated by a hybrid thermo-hydraulic model with regime uncertainty, and mapped to risk-aware choke/pump actions under probabilistic pressure-integrity constraints.

Another difficulty is that gas behaves as both a transported phase and a thermodynamic species that may partially dissolve into the liquid phase depending on pressure, temperature, and fluid chemistry. Dissolution changes both the volumetric signature of an influx and the rate at which it manifests as free gas. It also affects density and viscosity, thereby altering pressure losses and slip. The operational implications are direct: a dissolved fraction can delay or soften surface indicators, complicating early detection, while subsequent exsolution can accelerate volumetric growth during migration [2]. Accounting for such effects within a real-time framework requires thermodynamic consistency without introducing a full chemical process simulator into the loop. The framework in this paper addresses this by embedding a compact equilibrium module into the transport model, with a structure that supports fast evaluation and smooth sensitivities needed for filtering and control.

Finally, interpreting surface pressure signals requires recognizing that the same pressure trace can correspond to different internal states depending on whether the annulus is static, circulating, or transitioning through regime changes. In other words, the observation model is not stationary: the effective




Figure 2: Hybrid reduced-order state-space structure: a continuous thermo–two-phase state evolves under regime-conditioned closures, while a discrete latent regime follows a state-dependent Markov transition model that explains abrupt changes in slip and friction behavior.




Figure 3: Thermo–solubility coupling that reshapes the influx signature: local pressure/temperature determine equilibrium dissolved capacity, mass transfer redistributes gas between dissolved and free inventories, and transport with slip propagates the resulting compressibility and holdup effects into pressure evolution.

transfer function from downhole state to surface signals depends on flow regime, which itself may be a latent variable influenced by gas holdup, liquid rate, and geometry. This motivates a hybrid representation in which regime is modeled explicitly as a discrete state that modulates continuous dynamics. Such a representation makes it possible to perform inference over both the influx and the regime, and to quantify when the data are insufficient to discriminate among competing hypotheses.

This paper proposes a unified architecture that combines a mechanistic yet reduced thermo-hydraulic model, regime-switching closures, sequential Bayesian data assimilation, and risk-sensitive control [3]. The independent thesis is that real-time well-control decisions can be materially improved by treating kick detection and response as a single probabilistic digital twin, where the estimator and controller are co-designed around a hybrid reduced-order model that is explicitly constructed for identifiability from




Figure 4: Sequential Bayesian assimilation loop: prediction and measurement update maintain a hybrid posterior over influx and regime hypotheses, while parallel detection and detectability metrics quantify both event probability and how informative current telemetry is for identifying influx size and location.




Figure 5: Risk-sensitive decision layer: posterior uncertainty is propagated through a short-horizon optimizer that selects choke/pump actions while enforcing probabilistic pressure-integrity constraints and explicitly accounting for regime-dependent pressure-loss ambiguity.

surface telemetry. Rather than centering the contribution on any particular experimental dataset or simulator, the paper develops general mathematical structure: a conserved inventory formulation for free and dissolved gas, a regime-conditioned drift representation with bounded complexity, an information-theoretic detectability analysis guiding sensor utilization, and a chance-constrained control formulation that maps posterior uncertainty into pressure-integrity guarantees. The remainder of the paper derives the governing model, presents the hybrid reduction, develops the inference and control layers, and discusses computational guarantees and deployment considerations.




Figure 6: Information-aware closed loop: modest, constraint-compliant control perturbations can increase the information content of surface signals, accelerating posterior concentration on influx and regime hypotheses and improving safety margins without excessive conservatism.

Symbol	Quantity	Role in model	Units
ps, t	Pressure along annulus	Drives influx, slip, solubility and pressure-balance constraints	Pa
Ts, t	Temperature profile	Affects densities and gas K solubility c^*p, T	K
$\alpha s, t$	Free-gas holdup	Determines mixture density ρ_m and compressibility	–
$u_m s, t$	Mixture velocity	Governs convective transport of mass and energy	m/s
$c_d s, t$	Dissolved-gas concentration	Part of conserved inventory, controls hidden gas fraction	kg/m ³
$\rho_l p, T, c_d$	Liquid density	Enters hydrostatic and friction terms	kg/m ³
$\rho_g p, T$	Gas density	Links free-gas volume to mass; highly compressible	kg/m ³

Table 1: Key distributed state variables in the thermo-hydraulic influx model.

2. Physical and Information-Theoretic Model of Influx Evolution

Consider a wellbore annulus parameterized by measured depth coordinate $s \in [0, L]$ with $s = 0$ at the surface and $s = L$ at the bottomhole. Let t denote time. The annulus is represented as a one-dimensional conduit with cross-sectional area A_s and hydraulic diameter D_{hs} , potentially varying with geometry. The transported mixture consists of a liquid drilling fluid that may contain dissolved gas and a free-gas phase dispersed as bubbles, slugs, or other structures [4]. Let $\alpha s, t$ be the free-gas volumetric holdup, $u_m s, t$ the mixture velocity, ps, t the pressure, and Ts, t the temperature. Let $c_d s, t$ denote the dissolved-gas mass concentration in the liquid (mass of gas per unit liquid volume), and let $\rho_l p, T, c_d$ and $\rho_g p, T$ denote the liquid and gas densities, respectively. The mixture density is $\rho_m = 1 - \alpha \rho_l - \alpha \rho_g$. For operational relevance, the objective is not to resolve fine-scale interfacial dynamics but to represent the integrated impact of gas inventory, compressibility, and slip on the pressure field and on surface-observable quantities such as choke pressure, standpipe pressure shifts, and pit gain.

A convenient starting point is a conserved inventory formulation that separates free-gas mass from dissolved-gas mass while allowing interconversion through thermodynamic equilibrium or finite-rate transfer. Let m_{fs}, t be free-gas mass per unit length and m_{ds}, t be dissolved-gas mass per unit length.

Symbol	Definition	Role	Comment
m_{fs}, t	$As \alpha s, t \rho_g p, T$	Free-gas mass per unit length	Drives compressibility and slip response
m_{ds}, t	$As 1 - \alpha s, t c_d s, t$	Dissolved-gas mass per unit length	Can mask influx in early stages
m_{gs}, t	$m_f m_d$	Total gas mass per unit length	Conserved in absence of sources/sinks
$M_g t$	$\int_0^L m_{gs}, t ds$	Total gas inventory in well-bore	Primary influx size indicator
$\eta_d t$	$\frac{m_d ds}{m_g ds}$	Fraction of gas dissolved in liquid	Controls volumetric signature at surface
$\Gamma s, t$	Net transfer from dissolved to free gas	Couples m_f and m_d	Relaxation towards $c^* p, T$

Table 2: Gas inventory bookkeeping and phase-conversion variables in the conserved formulation.

Component	Description	Dimensionality	Notes
$p_s t$	Surface annulus pressure p_0, t	Scalar	Directly measured; boundary for profile
$p_b t$	Bottomhole pressure p_L, t	Scalar	Key for under/overbalance decisions
$a_i t$	Pressure-shape coefficients in p_s, t basis	n_p states	Capture hydrostatic and friction gradients
$b_i t$	Holdup-shape coefficients in $\alpha s, t$ basis	n_α states	Encode free-gas spatial distribution
$M_g t$	Total gas mass in annulus	Scalar	Integrates m_{gs}, t ; influx size summary
$\eta_d t$	Dissolved fraction of gas	Scalar in 0, 1	Links solubility to observed compressibility
$T_e t$	Effective temperature parameter	Scalar	Controls densities and $c^* p, T_e$

Table 3: Reduced-order continuous state vector used in the hybrid thermo-hydraulic model.

Using A as cross-sectional area, one may write

$$m_{fs}, t = As \alpha s, t \rho_g p s, t, T s, t, \quad m_{ds}, t = As 1 - \alpha s, t c_d s, t. \quad (2.1)$$

Total gas inventory per unit length is $m_g = m_f m_d$. When a kick enters, an initial condition injects gas mass into some interval near the bottom, either as free gas, dissolved gas, or both depending on thermodynamic conditions and mud chemistry. The subsequent evolution includes advection with the mixture, relative slip between phases, and redistribution between dissolved and free states as pressure and temperature vary.

For transport, a reduced two-fluid description can be expressed in terms of mixture continuity, gas inventory continuity, and mixture momentum, with a closure for slip. Let q_l be the liquid volumetric flow rate and q_g the free-gas volumetric flow rate. Define superficial velocities $j_l = q_l/A$ and $j_g = q_g/A$. The mixture velocity is $u_m = j_l + j_g$ [6]. A drift representation writes the gas phase velocity as $u_g = u_m + v_d \alpha, p, T; \theta$, where v_d is a drift or slip function with parameters θ capturing regime and rheology effects. The liquid phase velocity is $u_l = u_m - \frac{\alpha}{1-\alpha} v_d$ when only two phases are considered.

Regime r	Slip/friction behaviour	Typical conditions	Modeling focus
Dispersed-bubble	Small V_r , weak dependence of friction	Low gas rate, high liquid rate, nearly homogeneous flow	Smooth gradients, modest compressibility
Slug / intermittent	Moderate V_r , strong non-linearity vs. α	Intermediate gas fraction, transitional structures	Large pressure fluctuations, strong pit response
Churn / transitional	Highly variable V_r and losses	High gas fraction, complex structures	Uncertain closures; larger process noise
Annular-like	Large V_r , friction dominated by liquid film	Gas-dominated flow, liquid film on walls	Potentially high surface rates and pressure peaks

Table 4: Qualitative flow-regime classes encoded by the discrete latent variable rt and their closure implications.

Measurement	Symbol	Link to state	Main uncertainty contributors
Choke manifold pressure	p_{ct}	Function of p_{st} and surface hydraulics	Gauge bias, manifold dynamics
Surface annulus pressure	p_{st}	Boundary value of p_s, t	Sensor noise, wellbore storage
Return flow rate	$q_{out}t$	Depends on u_m, α at surface	Flow-meter bias, gas breakout at surface
Pit volume change	V_{pt}	Integral of volumetric expansion and storage	Level measurement error, thermal expansion
Pump rate / standpipe data	q_{int}	Known input, informs u_m	Actuator tracking, lag in reported set-points
Derived pressure features	ϕ_k	Functions of p_c, p_s over windows	Feature extraction error, aliasing

Table 5: Observation channels used in the digital twin and their relationship to the hybrid state.

In this representation, the free-gas holdup is determined implicitly by the relationship between j_g and u_g , namely $j_g = \alpha u_g$.

A compact gas inventory equation can be written by combining advection of free and dissolved components with a source term representing dissolution and exsolution. Let Γ, t be the net mass transfer rate from dissolved to free gas per unit length (positive for exsolution). Then

$$\frac{\partial m_f}{\partial t} \frac{\partial}{\partial s} (m_f u_g) = \Gamma, \quad \frac{\partial m_d}{\partial t} \frac{\partial}{\partial s} (m_d u_l) = -\Gamma. \quad (2.2)$$

This formulation conserves total gas mass in the absence of leakage and influx sources, while allowing redistribution between phases. For many operational windows, interphase mass transfer can be approximated as fast relative to advection so that the dissolved concentration tracks an equilibrium function of local pressure and temperature, $c_d \approx c^*p, T$, with a saturation limit. In other windows, finite-rate kinetics may matter, and Γ can be represented as a relaxation toward equilibrium, $\Gamma = \kappa A1 - \alpha (c_d - c^*p, T)$ with rate coefficient κ that may depend on turbulence and interfacial area. The digital-twin estimator can treat κ as uncertain and infer it if the telemetry is informative [7].

The equilibrium relation c^*p, T is the thermodynamic core that determines how much of the gas inventory can be hidden in the liquid. For synthetic-based or otherwise non-aqueous fluids, gas solubility can be substantial at downhole pressures, and neglecting dissolution can bias inferred free-gas volume and

Source of uncertainty	Model element	Probabilistic representation	Operational impact
Influx size and location	M_g , spatial holdup profile	Prior over M_g, b_i ; updated by filter	Drives kick severity and circulation time
Regime misclassification	Discrete rt	Markov chain with state-dependent transitions	Alters predicted pressure losses and slip
Solubility and fluid composition	c^*p, T , factor ζ	Random-walk with wide prior	Affects hidden gas fraction and detectability
Thermal environment	$T_e t$ and related parameters	Process noise on T_e , uncertain βs	Modulates density and expansion with depth
Model reduction error	fx, u vs. true dynamics	Additive process noise w_k	Limits estimator confidence and horizon length
Sensor bias and noise	y_k	Gaussian plus heavy-tailed components	Influences false alarms and state dispersion

Table 6: Principal uncertainty sources and how they are represented in the probabilistic digital twin.

Aspect	Implementation choice	Rationale	Expected benefit
State representation	Hybrid xt, rt with reduced bases	Capture essential physics with low dimension	Real-time feasibility with physical interpretability
Filtering architecture	Rao–Blackwellized particle filter	Sample discrete regime, linearize Gaussian part	Lower variance and better scalability
Constraint handling	Reparameterization (e.g. logistic for α)	Enforce physical bounds in state space	Avoid unphysical estimates under noise
Model comparison	Sliding-window likelihood ratio	Distinguish influx/no-influx hypotheses	Quantified detection confidence
Information metrics	Online Fisher/observability indicators	Monitor identifiability of M_g, s_c	Adjust control conservatism to data quality

Table 7: Summary of design choices in the Bayesian inference layer for influx detection.

thus predicted surface signatures. A mechanistic approach uses phase-equilibrium conditions expressed through equality of chemical potential or fugacity between phases. In a reduced implementation, the equilibrium module maps p, T to a dissolved fraction via an equation of state and activity or fugacity coefficients. A representative thermodynamic approach of this kind has been implemented in a simplified kick model by Manikonda et al. (2021), illustrating that an equilibrium-based solubility calculation can be integrated into fast kick simulation while retaining quantitative agreement with higher-fidelity process simulation [8]. In the present framework, the key requirement is not the particular equation of state but that the map c^*p, T be smooth, monotone in pressure over relevant ranges, and differentiable so that the filter can propagate sensitivities and the controller can assess risk margins.

The mixture momentum equation provides the link between gas inventory and pressure evolution, which is the primary observable at the surface [9]. In one-dimensional form, a reduced momentum balance for the mixture can be written as

$$\frac{\partial}{\partial t} (\rho_m u_m) - \frac{\partial}{\partial s} (\rho_m u_m^2) = -\frac{\partial p}{\partial s} - \rho_m g \cos \phi s - \frac{2fs, t}{D_h s} \frac{\rho_m u_m |u_m|}{2} \mathcal{M}_c, \quad (2.3)$$

Element	Mathematical form	Purpose	Example constraint or cost
Underbalance margin	$p_b t \geq p_{\text{pore}} \Delta_{\min}$	Prevent additional influx	Chance-constrained with high probability
Fracture integrity	$ps, t \leq p_{\text{frac}} s - \Delta_{\max}$	Avoid losses and formation damage	Enforced via tightened basis-coefficient bounds
Surface equipment limits	$p_c t \leq p_{\max}$	Protect choke/manifold hardware	Hard limit in optimization
Actuation limits	$\Delta u t$ bounded	Respect choke and pump rateability	Rate and amplitude constraints in MPC
Risk measure	CVaR $_{\alpha}$ Zu of cost Z	Emphasize tail events over mean cost	Penalize scenarios with high peak pressure
Dual-control term	Information penalty $\Sigma_{kH} u$	Trade off exploration vs control	Prefer inputs that reduce state uncertainty

Table 8: Risk-sensitive control objectives and constraints for closed-loop well-control synthesis.

Limitation	Impact on predictions	Mitigation in framework	Possible extension
One-dimensional annulus	Misses azimuthal asymmetry and beds	Inflate process noise, regime-aware closures	Two-compartment (high/low side) model
Simplified rheology	Approximate Newtonian effects	non-Effective viscosity and friction multipliers	Explicit cuttings and rheology states
Coarse thermal model	Under-resolved temperature transients	Treat T_e as uncertain, adapt from data	Low-order temperature profile basis
Equilibrium-biased solubility	May misrepresent finite-rate transfer	Uncertain factor ζ and rate κ	Richer mass-transfer kinetics, calibration
Limited sensing	Surface-only observability gaps	Detectability metrics and conservative control	Integrate downhole and distributed sensing

Table 9: Selected modeling limitations, their consequences, and directions for refinement within the digital-twin architecture.

where g is gravitational acceleration, ϕs is inclination angle from vertical, f is a friction factor generalized to non-Newtonian effective viscosity, and $\mathcal{M}_c$ denotes correction terms capturing acceleration, cuttings loading, and other effects that may be lumped into uncertain parameters if not directly modeled. The friction factor and slip closure are precisely where regime dependence enters: as gas holdup increases and structures transition from dispersed bubbles to intermittent or churn-like behavior, both slip and pressure losses can change nonlinearly.

Temperature affects both densities and solubility. A reduced thermal model can be introduced to capture the dominant effect: along-well heat exchange with formation and viscous dissipation. Denote by h_{wfs} an effective heat-transfer coefficient from wellbore fluid to formation and by T_{fs} formation temperature profile. A mixture energy balance, simplified to a convective-relaxation form, can be written as

$$\frac{\partial T}{\partial t} u_m \frac{\partial T}{\partial s} = \beta s (T_{fs} - T) \chi s u_m^2, \quad (2.4)$$

where β aggregates heat transfer per unit thermal mass and χ aggregates viscous heating effects. In many real-time contexts, the energy model is used primarily to correct density and solubility rather than to predict detailed temperature transients, so β and χ can be coarsely parameterized and inferred if needed [10].

The purpose of formalizing these equations is not to claim a universally accurate multiphase simulator, but to define a state-space model from which one can quantify observability of the gas inventory from surface telemetry. Let y_t represent measured signals such as choke pressure p_{ct} , pump rate, return flow rate, and pit volume. These are functions of the distributed state through boundary conditions and integrals. For example, choke pressure is linked to p_0, t through surface manifold dynamics, while pit gain is linked to the integral of volumetric expansion of the mixture relative to the baseline. A reduced observation model can be written as $y = \mathcal{H}p, \alpha, c_d, T \nu$, where ν is measurement noise.

To guide reduction, it is useful to consider the Fisher information content of y with respect to parameters of interest such as total gas mass $M_g = \int_0^L m_g ds$, gas centroid location s_{ct} , and solubility scaling. In a linearized setting around a nominal trajectory, let ϑ denote a vector of such quantities and $y_t; \vartheta$ the predicted measurement. Under Gaussian noise with covariance Rt , the Fisher information accumulated over a window t_0, t_1 is

$$\mathcal{I}\vartheta = \int_{t_0}^{t_1} \left(\frac{\partial y}{\partial \vartheta} \right)^\top R^{-1} \left(\frac{\partial y}{\partial \vartheta} \right) dt. \quad (2.5)$$

A reduced-order model is valuable if it preserves the dominant sensitivities $\partial y \partial \vartheta$ that determine detectability [11]. This criterion motivates retaining explicit states for gas inventory and its conversion between dissolved and free forms, because these directly influence compressibility and thus pressure response. It also motivates introducing a discrete regime variable because it can cause discontinuous changes in sensitivities that would otherwise appear as unexplained noise.

3. Reduced-Order Two-Phase Thermo-Hydraulics With Regime Switching

A distributed model is difficult to embed directly in a real-time estimator and controller, particularly when the inference must explore multiple hypotheses over influx size and location. Reduction is therefore necessary, but it should be structured to preserve conservation, monotonicity, and identifiable mappings to surface data. The approach here is to represent the state with a small number of physically interpretable lumped variables that approximate the pressure profile and gas distribution, while representing regime effects through a latent discrete variable that modulates closure parameters.

Let the continuous state vector be x_t , containing bottomhole pressure $p_{bt} = p(L, t)$, surface pressure $p_{st} = p(0, t)$, a small set of coefficients $a_i t$ describing the pressure profile along s , a set of coefficients $b_i t$ describing the free-gas holdup profile, total gas mass $M_g t$, dissolved fraction $\eta_d t = \frac{m_d ds}{m_g ds}$, and an effective temperature parameter $T_e t$ capturing average thermodynamic conditions relevant for solubility. The pressure profile can be approximated by a basis expansion

$$p(s, t) \approx p_{st} + \sum_{i=1}^{n_p} a_i t \psi_i(s), \quad (3.1)$$

where ψ_i are smooth basis functions satisfying $\psi_i(0) = 0$ and scaled so that $\int_0^L a_i \psi_i ds = p_b - p_s$ [12]. Similarly, the holdup profile can be approximated as

$$\alpha(s, t) \approx \sum_{i=1}^{n_\alpha} b_i t \varphi_i(s), \quad (3.2)$$

with φ_i chosen nonnegative and bounded so that the approximation respects $\alpha \in [0, 1]$. The gas inventory M_g then constrains the combination of α and c_d through the density relation. This representation turns the PDE model into a set of coupled ordinary differential equations for the coefficients, obtained by Galerkin projection or by enforcing integral balances on segments.

The key nonlinearity is how gas inventory affects pressure gradients through compressibility, density, and friction. In the reduced model, this effect is captured by defining an effective hydrostatic-plus-friction map that links p_b and p_s given x and controls u (such as pump rate and choke opening). A generic reduced pressure-drop relation can be written as

$$p_b - p_s = \int_0^L \rho_m s, t g \cos \phi s ds + \int_0^L \Delta p_f s, t ds + \Delta p_a t, \quad (3.3)$$

where Δp_f represents frictional gradient and Δp_a represents acceleration and minor-loss contributions. In a basis representation, these integrals become functions of the coefficients a_i , b_i , T_e , and closure parameters.

Flow regime affects both slip and friction. Rather than attempting to deterministically map state to regime via hard thresholds, the digital twin models regime as a discrete latent variable rt taking values in a finite set representing broad structural classes of two-phase behavior. Each regime defines a parameter vector θ_r controlling drift and friction closures. The continuous dynamics are then conditioned on r , creating a hybrid system:

$$xt = f(xt, ut; \theta_{rt}) \quad wt, \quad (3.4)$$

$$\Pr(rt \Delta t = j \mid rt = i, xt) = \Pi_{ij}(xt), \quad (3.5)$$

where wt aggregates model error and unresolved dynamics, and Π_{ij} is a transition probability matrix that can depend smoothly on gas holdup and flow rate, reflecting that certain transitions are more likely as conditions change. The discrete-state formulation provides two advantages [14]. It allows the estimator to attribute abrupt changes in pressure response to plausible regime transitions rather than to spurious influx changes, and it allows the controller to hedge decisions against uncertainty in regime-dependent pressure losses.

A practical reduced closure can be defined by writing the slip velocity as

$$v_d = V_r \mathcal{S}(\alpha; \lambda_r) \mathcal{G}(p, T_e), \quad (3.6)$$

where V_r is a regime-dependent scale, $\mathcal{S}$ is a bounded shape function of holdup with parameter λ_r , and $\mathcal{G}$ captures compressibility and density ratio effects. The friction gradient can be written as

$$\Delta p_f = K_r \mu_{\text{eff}}^{m_r} \rho_m^{n_r} |u_m|^{\ell_r}, \quad (3.7)$$

where μ_{eff} is an effective viscosity computed from rheology and holdup, and exponents m_r, n_r, ℓ_r encode regime dependence while maintaining dimensional consistency via K_r . Such a parametric form is not intended as a universal correlation; rather, it is chosen to be expressive enough to capture dominant trends while having a small number of parameters that can be inferred online.

Static and circulating conditions differ in how pressure signals propagate and in which regimes are prevalent. Under static shut-in conditions, the absence of imposed liquid circulation can lead to different holdup structures and pressure-gradient behavior than under dynamic circulation, even for similar gas inventories. Experimental observations linking annular flow patterns to measured pressure gradients under static conditions have been explored using machine learning classification approaches, showing that surface pressure responses can be informative of flow structure when mapped through an appropriate model [15]. Within the present hybrid formulation, such empirical insights are used only to motivate the explicit inclusion of a regime variable and to justify using surface pressure features as informative observations for rt ; the digital twin itself remains grounded in conservation and pressure-balance structure [16].

A reduced observation model must map the hybrid state to surface measurements. Let yt include surface pressure p_s , choke manifold pressure p_c , return flow rate q_{out} , and pit volume change V_p . A

compact mapping can be written as

$$yt = h(xt, ut; \theta_{rt}) \quad vt, \quad (3.8)$$

where vt is measurement noise. For example, pit gain is related to the net volumetric expansion of the mixture, which depends on α and on dissolved fraction through exsolution. In the reduced model, one can express the incremental pit volume rate as the difference between outlet and inlet volumetric flow rates, corrected for compressibility and storage in the wellbore:

$$V_p t \approx q_{out} t - q_{in} t \frac{L}{0} \frac{\partial}{\partial t} (As \alpha s, t) \quad ds \quad Ct, \quad (3.9)$$

where Ct accounts for elastic wellbore storage and thermal expansion of the liquid. Because α and its time derivative depend strongly on solubility and pressure evolution, pit gain becomes a coupled observable that can help disambiguate whether a pressure anomaly corresponds to dissolved gas slowly evolving versus a free-gas slug rapidly migrating.

Reduction introduces approximation error, so it is essential to bound its effect on inference. Let x^*t be the state of a higher-fidelity model and xt the reduced state embedded in the estimator [17]. The mismatch can be represented as a disturbance wt whose covariance or bounds are calibrated to ensure that posterior uncertainty does not become falsely confident. A useful approach is to choose wt so that the reduced model is conservative in the sense of enclosing high-fidelity trajectories within a tube with high probability, which can be formalized through incremental stability and Lipschitz properties discussed later. The hybrid structure also ensures that certain classes of mismatch, such as sudden changes in slip behavior, can be absorbed by transitions in rt rather than corrupting the continuous estimate.

4. Bayesian Data Assimilation From Surface Measurements

Given the reduced hybrid model, the inference task is to compute a posterior distribution over the continuous state xt and regime rt conditioned on the measurement history $\mathcal{Y}_t = \{y\tau : \tau \leq t\}$. The posterior must be updated online at operational sampling rates, and it must remain stable under non-Gaussianity induced by regime switching, constraints $\alpha \in [0, 1]$, and nonlinear solubility effects. A sequential Bayesian approach provides a principled mechanism for propagating uncertainty, detecting when data do not support a unique inference, and producing confidence intervals needed for risk-aware control.

Let the system be discretized at times t_k with step Δt . The hybrid dynamics can be written as [18]

$$x_{k+1} = f(x_k, u_k; \theta_{r_k}) \quad w_k, \quad (4.1)$$

$$r_{k+1} \sim \Pi(\cdot \mid r_k, x_k), \quad (4.2)$$

$$y_k = h(x_k, u_k; \theta_{r_k}) \quad v_k, \quad (4.3)$$

with process noise w_k and measurement noise v_k . The goal is $px_k, r_k \mid \mathcal{Y}_k$. Because r_k is discrete and x_k is continuous and constrained, exact filtering is generally intractable, motivating approximate methods.

A suitable approach is a marginalized particle filter in which particles represent hypotheses over the discrete regime trajectory and selected nonlinear continuous components, while conditionally Gaussian components are updated by Kalman-like recursions. To enable marginalization, the continuous state can be partitioned as $x = x^n, x^g$, where x^n contains strongly nonlinear variables such as dissolved fraction and holdup coefficients, and x^g contains variables that can be linearized locally such as pressure basis coefficients. Conditional on x^n, r , the mapping from x^g to y may be approximated as locally affine with Jacobian computed from the reduced model. This yields a Rao–Blackwellized structure that reduces variance compared to a fully sampled particle filter.

Let particle i at time k carry $x_k^{n,i}, r_k^i$ and an associated conditional Gaussian distribution over x_k^g with mean μ_k^i and covariance P_k^i . The prediction step propagates each particle through the nonlinear dynamics and regime transition:

$$x_{k1|k}^{n,i} = f^n \left(x_k^{n,i}, \mu_k^i, u_k; \theta_{r_k^i} \right) w_k^{n,i}, \quad (4.4)$$

$$r_{k1}^i \sim \Pi \left(\cdot \mid r_k^i, x_k^{n,i} \right), \quad (4.5)$$

while the Gaussian part is predicted by linearization around the particle:

$$\mu_{k1|k}^i = f^g \left(x_k^{n,i}, \mu_k^i, u_k; \theta_{r_k^i} \right), \quad (4.6)$$

$$P_{k1|k}^i = A_k^i P_k^i A_k^{i\top} Q_k^i, \quad (4.7)$$

where A_k^i is the Jacobian of f^g with respect to x^g and Q_k^i is process noise covariance. The update step uses the measurement model linearized similarly: [19]

$$y_k \approx h \left(x_{k|k-1}^{n,i}, \mu_{k|k-1}^i, u_k; \theta_{r_k^i} \right) H_k^i \left(x_k^g - \mu_{k|k-1}^i \right) v_k, \quad (4.8)$$

leading to a Kalman update for μ_k^i, P_k^i and an importance weight update for the particle based on the innovation likelihood.

This formulation naturally incorporates constraints by reparameterization. For example, holdup coefficients can be mapped through a logistic function to enforce $\alpha \in [0, 1]$, and dissolved fraction $\eta_d \in [0, 1]$ can be similarly constrained. The solubility equilibrium map c^*p, T_e introduces strong nonlinear coupling between pressure and dissolved inventory. Within the filter, it is beneficial to treat the equilibrium map as deterministic given p, T_e but to allow a multiplicative uncertainty factor ζ capturing fluid-composition variability and model mismatch, i.e., $c^* = \zeta \hat{c}^*p, T_e$. The factor ζ can be included as a slowly varying state with a random-walk model, enabling the filter to adapt when the true solubility deviates from nominal due to mud formulation differences.

Detecting an influx is not merely estimating its magnitude; it is also deciding whether the posterior probability that an influx exists exceeds a threshold. A statistically consistent approach introduces a model class $\mathcal{M}_0$ representing no influx and $\mathcal{M}_1$ representing influx, each with its own state evolution. The filter then computes $p\mathcal{M}_1 \mid \mathcal{Y}_k$ via Bayesian model selection with an evidence ratio. In practice, to avoid fragile dependence on priors, one can compute an approximate log-likelihood ratio over a sliding window:

$$\Lambda_k = \log_{j=k-W+1}^k \frac{py_j \mid \mathcal{Y}_{j-1}, \mathcal{M}_1}{py_j \mid \mathcal{Y}_{j-1}, \mathcal{M}_0}, \quad (4.9)$$

and trigger detection when Λ_k exceeds a calibrated threshold that accounts for false-alarm tolerance and operational context [20]. Because the hybrid regime state can explain certain transients without invoking an influx, including r in both models reduces false positives that would otherwise arise from regime-induced pressure changes.

An important aspect of this inference problem is that some variables are only weakly identifiable from surface measurements at certain times. For example, early after influx entry, dissolved gas may dominate and free-gas holdup may be small, yielding limited pit gain while pressure changes can be subtle. To quantify this, the filter can compute an online observability score based on empirical Fisher information approximations derived from particle sensitivities. For each particle, one can propagate a sensitivity matrix $S_k = \partial \hat{y}_k / \partial \vartheta$ with respect to parameters ϑ such as M_g and s_c , and aggregate across particles. The smallest singular value of an information matrix approximation then indicates when the data are insufficiently informative, which is critical for safe decision making: if identifiability is low, the controller must act conservatively even if the mean estimate appears benign.

The hybrid filter can also incorporate feature-level information without sacrificing physical structure. Surface pressure signals contain transient signatures such as oscillations and slope changes that correlate with regime and with gas migration [21]. Rather than feeding raw features into an opaque predictor, the framework uses them as additional noisy observations of the regime state r_k through an auxiliary likelihood $p\phi_k \mid r_k$, where ϕ_k denotes extracted features such as normalized pressure derivative or spectral content over a window. This allows data-driven signal interpretation to inform the discrete regime estimate while the continuous physical state remains governed by conservation and pressure balance. The result is a coherent posterior that merges mechanistic and statistical evidence in a single probabilistic object.

Computationally, the filter must operate under strict latency constraints. The reduced model supports this by limiting state dimension and by enabling parallel particle propagation. Further reduction can be achieved by adaptively varying particle count based on posterior dispersion: when the posterior becomes concentrated, fewer particles suffice; when regime ambiguity increases, more particles are allocated. Resampling is performed cautiously to avoid particle impoverishment, especially in parameters like ζ that evolve slowly. Stability is enhanced by injecting small rejuvenation noise and by enforcing physical consistency checks such as nonnegative gas inventory and bounded holdup [22]. These measures are not mere implementation details; they are necessary to ensure that the digital twin remains well behaved in the face of real telemetry, which can contain dropouts, delays, and sensor biases.

5. Risk-Sensitive Well-Control Synthesis Under Posterior Uncertainty

Estimating the influx state is only useful insofar as it informs a control policy that reduces risk while respecting operational constraints. In well control, actions include adjusting choke opening, modifying pump rate, and selecting circulation strategies. These actions influence the pressure profile and thus the driving forces for further influx or for safe removal. The key difficulty is that the state is uncertain and the system is nonlinear and hybrid. A deterministic controller designed for a point estimate can lead to unsafe decisions if the posterior has heavy tails or if regime ambiguity causes underestimation of pressure losses. Therefore, the control problem is formulated as risk-sensitive optimization conditioned on the full posterior produced by the estimator [23].

Let ut denote the control input vector, such as choke coefficient and pump rate. Let xt and rt denote the hybrid state. Define a constraint function $g(x, u, r)$ representing pressure integrity margins. Typical constraints include keeping bottomhole pressure above pore pressure by a minimum margin to prevent additional influx, and keeping wellbore pressure below fracture pressure by a margin to prevent losses. Additional constraints include surface equipment pressure limits and actuation rate limits. Because x and r are uncertain, constraints should hold with high probability. A chance constraint can be written as [24]

$$\Pr(g(xt, ut, rt) \leq 0 \mid \mathcal{Y}_t) \geq 1 - \epsilon, \quad (5.1)$$

where ϵ is a small risk tolerance. The function g can include multiple components; for example, one component may be $p_{bt} - p_{\text{pore}}L \geq \Delta_{\min}$ and another may be $p_{s,t} - p_{\text{frac}}s \leq -\Delta_{\max}$ for all s , approximated in the reduced model by constraints on basis coefficients.

A tractable formulation over a finite horizon t, tH is a receding-horizon risk-sensitive optimal control problem. Let J be a cost functional penalizing deviation from desired pressures, excessive control effort, and time to circulate out the gas inventory. Conditioning on the posterior, a risk-sensitive objective can be written using an exponential utility or a coherent risk measure. A practical choice is conditional value-at-risk (CVaR), which focuses on tail outcomes. For a random cost Zu induced by uncertainty in x, r and process noise, CVaR at level α is

$$\text{CVaR}_\alpha Z = \min_{\tau \in \mathbb{R}} \left(\tau + \frac{1}{1 - \alpha} \mathbb{E}[Z - \tau] \right). \quad (5.2)$$

The control problem then becomes minimizing $\text{CVaR}_\alpha Z u$ subject to chance constraints. In a particle-based implementation, the expectation is approximated over posterior samples, and chance constraints can be enforced by requiring that constraints hold for a specified fraction of particles, possibly with tightening to account for sampling error.

The hybrid regime state introduces additional complexity because different regimes yield different predicted pressure drops for the same u [25]. To hedge against regime ambiguity, the controller can optimize for a mixture model. If the posterior assigns probabilities π_j to regimes $r = j$, then predicted outcomes are mixtures, and constraints can be required to hold for each regime with weighted risk. Alternatively, one can enforce robust constraints against the worst-case regime among those with non-negligible posterior mass:

$$\max_{j \in \mathcal{R}_\delta} \Pr(gx, u, j \leq 0 \mid \mathcal{Y}_t, r = j) \geq 1 - \epsilon, \quad (5.3)$$

where $\mathcal{R}_\delta$ includes regimes with posterior probability above δ . This approach avoids catastrophic underestimation of friction when the system might transition into a higher-loss regime.

Control actions also influence observability. For example, adjusting choke or pump rate can create excitation that makes pressure signals more informative about gas inventory and regime. This is an instance of dual control: actions affect both the physical state and the information gained [26]. A digital twin can incorporate this by adding an information gain term to the objective, such as maximizing expected Fisher information or minimizing posterior entropy over M_g and s_c . Let $\Sigma_{kH} u$ denote a predicted posterior covariance after applying a candidate control sequence. An information-aware objective could include $\log \det \Sigma_{kH} u$ as a penalty. Because exact prediction of posterior under future measurements is expensive, an approximation can be used in which the filter is propagated forward with assumed measurement noise and linearized dynamics, yielding a predicted information matrix. Even approximate information-aware control can reduce uncertainty more quickly, leading to safer and less conservative decisions.

The control layer must also handle the thermodynamic coupling that can cause delayed exsolution. If a significant portion of gas is dissolved at depth, maintaining higher bottomhole pressure can keep gas in solution longer, potentially reducing surface gas handling demands but prolonging the period of uncertainty. Conversely, reducing pressure to expedite circulation can cause exsolution and increase free-gas volume, raising surface pressure and pit gain rates. The risk-sensitive formulation naturally captures this trade-off because dissolved fraction is part of the state distribution and because constraints are expressed probabilistically. The controller can therefore choose a pressure trajectory that balances integrity margins against operational efficiency, without assuming that the dissolved-to-free conversion is negligible [27].

To ensure safe closed-loop behavior, it is useful to establish sufficient conditions for constraint satisfaction under bounded disturbances. Suppose the reduced dynamics are incrementally input-to-state stable with respect to disturbances w within a set $\mathcal{W}$, and suppose the estimator provides a credible set $\mathcal{X}_t$ containing the true state with probability at least $1 - \delta$. If the controller ensures that constraints hold for all $x \in \mathcal{X}_t$ and all $w \in \mathcal{W}$ over the horizon, then constraints will hold in closed loop with probability at least $1 - \delta$ minus approximation errors. While exact robust tube MPC can be conservative, this reasoning motivates the practical strategy of combining posterior credible intervals with model-error bounds to tighten constraints. The hybrid regime variable complicates these guarantees, but by maintaining separate credible sets conditioned on regime and enforcing constraints across the mixture, one can obtain a conservative yet computationally feasible safety envelope.

6. Computational Architecture, Stability, and Error Bounds

A real-time digital twin must operate reliably under constraints that are not present in offline modeling: limited computation, intermittent telemetry, sensor biases, and the need for transparent uncertainty reporting. The architecture implied by the preceding sections consists of a reduced hybrid simulator, a

sequential Bayesian estimator, and a risk-sensitive controller operating in a receding-horizon loop. This section discusses how to structure computation, how to maintain numerical stability, and how to derive error bounds that support safe decision making.

The reduced model must be evaluated thousands of times per minute when embedded in a particle-based filter and an optimization-based controller [28]. To enable this, the model should have a small state dimension, avoid stiff dynamics, and provide analytic or efficiently computed Jacobians. The equilibrium solubility map should be implemented as a smooth function with bounded derivatives across expected ranges of p, T_e . If an equation of state is used, it should be reduced to closed-form or tabulated approximations to avoid iterative root finding inside the filter. Consistency constraints, such as conservation of total gas mass absent sources, should be enforced numerically through correction steps that project states back onto physically admissible manifolds. For instance, after each propagation step, holdup coefficients can be adjusted to ensure $\alpha_s \in [0, 1]$ and that the implied free-gas mass does not exceed total gas inventory.

Estimator stability depends on avoiding overconfidence. A common failure mode is that the filter underestimates model error and thus collapses posterior variance, producing brittle control decisions [29]. This is mitigated by explicit process noise w_k that reflects reduction error and closure uncertainty. One can adapt Q_k online by monitoring normalized innovation squared statistics. If innovations are consistently larger than expected under the assumed measurement noise, the process noise is inflated. Conversely, if innovations are consistently smaller, noise can be reduced but never below a floor representing irreducible mismatch. Such adaptive noise tuning can be justified as an empirical Bayes procedure that maximizes marginal likelihood, but in deployment it primarily serves to keep uncertainty calibrated.

A hybrid system introduces additional numerical challenges because regime transitions can cause discontinuities in predicted measurements. If the filter treats regimes as hard switches, particles can receive near-zero likelihood and degeneracy can occur. This is mitigated by smoothing the dependence of closures on r through regime-conditioned parameters that vary continuously within each regime and by allowing soft transitions through $\Pi_{ij}x$ that depend on continuous state. In addition, the observation likelihood can incorporate a small heavy-tailed component to prevent outlier measurements from annihilating particles, reflecting the reality that surface pressure can exhibit spikes due to equipment transients [30]. The heavy-tailed likelihood can be implemented as a mixture of Gaussians with a small weight on a large-variance component, which retains robustness without abandoning probabilistic semantics.

Beyond numerical stability, it is useful to establish analytical bounds that connect reduction error to estimation and control error. Consider the reduced continuous dynamics $\dot{x} = f(x, u; \theta_r)$ and suppose f is locally Lipschitz in x uniformly over admissible u and over regimes, with Lipschitz constant L_f . Suppose the true dynamics differ by an additive disturbance dt with $\|dt\| \leq \bar{d}$. Then the state error $et = x^*t - xt$ between true and reduced trajectories under the same input satisfies the Grönwall bound

$$\|et\| \leq \|e0\| e^{L_f t} \int_0^t e^{L_f t - \tau} \bar{d} d\tau \quad \Rightarrow \quad \|et\| \leq \|e0\| e^{L_f t} \frac{\bar{d}}{L_f} (e^{L_f t} - 1), \quad (6.1)$$

provided $L_f > 0$. This bound is conservative but useful: it motivates keeping L_f small through appropriate state parameterization and scaling, and it motivates limiting horizon length in the controller when model mismatch is high. In the hybrid case, one can apply such bounds piecewise between regime transitions and then incorporate additional jump errors due to parameter changes, which can be bounded if regime-conditioned parameters lie within known sets [31].

For the estimator, if the observation model is locally Lipschitz and the measurement noise is bounded in variance, then posterior covariance cannot shrink arbitrarily fast unless the system is strongly observable. The online information score mentioned earlier provides a practical detectability indicator. When detectability is low, the estimator should report wide credible intervals and the controller should enforce more conservative constraints. This coupling between detectability and control conservatism is a core

principle of safe probabilistic decision making: confidence should be earned from informative data rather than assumed.

The computational architecture should also handle telemetry delays. Surface measurements may arrive with latency, and some measurements such as return flow can be delayed relative to downhole events. This can be handled by augmenting the state with delay lines or by applying a fixed-lag smoother that reprocesses the last W measurements [32]. A fixed-lag smoother can substantially improve regime inference because regime transitions may be better inferred after observing subsequent pressure evolution. Implementation-wise, smoothing can be performed by storing particle ancestry over the lag window and reweighting. The control layer can operate on the filtered posterior for immediacy while the smoother refines historical estimates for diagnostics and model adaptation.

Validation of such a digital twin is not solely about matching a particular trajectory but about calibrating uncertainty. The appropriate evaluation asks whether true states lie within predicted credible intervals at the claimed rates, whether detection triggers have controlled false-alarm probabilities, and whether control actions maintain constraints under uncertainty. When laboratory or field datasets are available, one can use posterior predictive checks: simulate from the posterior over states and parameters and compare simulated measurement distributions to observed measurements, focusing on features relevant for decisions such as maximum surface pressure and pit gain rates. If the digital twin systematically underpredicts extremes, risk constraints should be tightened.

The regime-switching concept is particularly relevant under dynamic circulation, where the relationship between pressure gradient and flow structure can shift with gas rate and viscosity [33]. Experimental work under dynamic conditions has shown that supervised and unsupervised learning methods can classify flow regimes and relate them to surface pressure and flow data with high accuracy [34]. In the present framework, the significance of such results is that they support the feasibility of treating regime as an inferable latent variable from surface data, enabling the hybrid filter to allocate probability mass across regimes rather than collapsing to a single assumed closure. The digital twin does not require the same classification models; it uses the broader idea that regime leaves a detectable footprint in surface telemetry, which can be integrated into a probabilistic hybrid estimator.

7. Discussion of Limitations and Extensions

Any reduced probabilistic model for well control must navigate a trade space among physical fidelity, identifiability, and computational tractability. The framework developed here is designed to be deployable, but it necessarily abstracts away certain effects that can matter in specific contexts. Discussing these limitations is not a caveat to the approach but a practical guide to where model extension is warranted and how uncertainty should be represented when extensions are not feasible.

One limitation is the one-dimensional assumption. Annular flow can exhibit azimuthal asymmetries, especially in deviated wells, eccentric annuli, and in the presence of cuttings beds [35]. Such effects can alter holdup distribution and effective friction. In the reduced model, these are lumped into regime-dependent parameters and process noise. This is reasonable if the surface telemetry cannot uniquely identify azimuthal structure, but it may become insufficient when severe eccentricity causes systematic bias in pressure gradients. A possible extension is to use a two-compartment model representing high-side and low-side flow paths, coupled by crossflow, still maintaining low dimensionality. The hybrid regime variable can then represent not only vertical flow patterns but also bed formation states, again treated probabilistically.

Another limitation concerns cuttings transport and non-Newtonian rheology. Drilling fluids often exhibit yield stress and shear-thinning, and cuttings concentration can change effective density and viscosity. The reduced friction parameterization includes an effective viscosity term, but in deployment one may need to explicitly track cuttings concentration as a state if it varies rapidly or if it covaries with influx signatures [36]. The challenge is that cuttings concentration may be weakly observable from surface pressure alone. Incorporating it without additional measurements can inflate uncertainty and complicate

inference. A practical compromise is to represent cuttings effects as a slowly varying friction multiplier inferred over longer windows, thereby capturing drift without overfitting short-term transients.

Thermal modeling is another area where reduction can both help and hinder. Temperature affects solubility and density, but surface temperature measurements may be limited and downhole temperature may be unobserved. The reduced thermal parameter T_e is intended to capture the main effect on solubility. However, if significant thermal transients occur, such as during prolonged circulation changes, a single parameter may be inadequate [37]. Extending the thermal model to a low-order profile expansion similar to the pressure basis can improve fidelity, but it also increases state dimension and may degrade identifiability. In such cases, the digital twin should treat thermal uncertainty explicitly and propagate it into solubility uncertainty, thereby avoiding false confidence in dissolved fraction estimates.

The equilibrium dissolution assumption can also be context dependent. In some cases, mass transfer between free and dissolved gas may not be instantaneous, especially if flow is laminar or if interfacial area is limited by large gas structures. The relaxation model with rate κ partially addresses this, but inferring κ online may be difficult because its effect can be confounded with regime and with gas inventory. Here, identifiability analysis is critical: if Fisher information indicates that κ is not identifiable from available signals, the estimator should not attempt to estimate it aggressively. Instead, it can marginalize over a plausible range of κ , leading to wider but more honest uncertainty in predicted pit gain and pressure evolution [38].

The controller formulation emphasizes chance constraints and tail risk, but operational decision making also depends on procedural constraints and human factors that are difficult to encode mathematically. For example, certain choke adjustments may be disallowed by procedures, or operators may require interpretability and stable recommendations rather than rapidly changing controls. These considerations can be incorporated by adding penalties on control rate changes and by constraining the controller to a subset of admissible policies. One can also implement a supervisory layer that translates the continuous optimal control outputs into discrete procedural recommendations, while retaining the posterior risk metrics to inform operator judgment.

A key extension is to integrate multi-source data. While the paper focuses on surface pressure and flow signals, modern operations may include downhole pressure-while-drilling, acoustic measurements, or distributed fiber sensing. Each additional sensor can dramatically improve observability and reduce posterior uncertainty. The framework is designed to accommodate this by expanding the observation model hx, u and measurement vector y [39]. The main requirement is that each sensor be characterized by an uncertainty model and latency model so that the filter can correctly weight information. Adding downhole pressure transforms the problem: the pressure profile becomes partially observed, and the estimator can disentangle friction changes from gas inventory changes more robustly. In such settings, the hybrid regime variable remains valuable because regime affects how pressure differences between sensors map to holdup and slip.

Finally, the relationship between the reduced model and any data-driven components should be treated carefully. The hybrid filter can incorporate data-driven regime likelihoods or feature models, but these must be calibrated to avoid injecting overconfident classifications that override physical consistency. A principled way to do this is to treat data-driven outputs as soft evidence with explicit uncertainty, and to allow the filter to disagree when physical constraints make a classification implausible. This preserves the integrity of the probabilistic semantics and prevents a learned classifier from becoming a single point of failure [40].

8. Conclusion

This paper presented a real-time probabilistic digital-twin framework for gas influx detection and well-control decision making that unifies reduced thermo-hydraulic modeling, thermodynamic dissolution effects, regime-switching closures, sequential Bayesian inference, and risk-sensitive control. The core methodological contribution is a hybrid reduced-order state-space model that conserves gas inventory while explicitly representing conversion between dissolved and free gas, and that accounts for

regime-dependent slip and friction through a discrete latent variable. Building on this model, the paper developed an online estimator that assimilates surface telemetry to produce posterior distributions over influx size, location, and key transport parameters, accompanied by detectability metrics that indicate when the data support confident inference. A risk-aware control layer was then formulated using chance constraints and tail-risk objectives to map posterior uncertainty into safe choke and circulation policies under integrity constraints.

The overall thesis is that well-control decisions benefit from treating detection and response as a single probabilistic closed-loop problem rather than as separate deterministic steps. By explicitly propagating model and regime uncertainty, the framework avoids brittle reliance on point estimates and provides a structured basis for conservative action only when warranted by limited information. The reduced formulation targets field deployment by emphasizing computational tractability, calibrated uncertainty, and bounded-error reasoning. Future work within this framework includes extending the reduced state to capture eccentric annulus effects and cuttings dynamics when telemetry supports it, integrating additional sensing modalities to strengthen observability, and refining model-error representations to improve calibration across diverse well geometries and mud systems [41].

References

- [1] A. T. Bozdana, N. Al-Kharkhi, and K. Al-Kharkhi, “Comparative experimental and numerical investigation on electrical discharge drilling of aisi 304 using circular and elliptical electrodes,” *Strojniški vestnik - Journal of Mechanical Engineering*, 4 2018.
- [2] R. A. Leishear, “A hydrogen ignition mechanism for explosions in nuclear facility piping systems,” *Journal of Pressure Vessel Technology*, vol. 135, pp. 054501–, 9 2013.
- [3] S. Wolhart, “Stimulation technologies for deep well completions,” 6 2003.
- [4] P. Shaw, D. Nickelson, and R. Hyde, “Operable unit 7-13/14 in situ thermal desorption treatability study work plan,” 5 1999.
- [5] C. P. Gooneratne, B. Li, and M. T. E, “Downhole applications of magnetic sensors,” *Sensors (Basel, Switzerland)*, vol. 17, pp. 2384–, 10 2017.
- [6] L. F. de Campos, P. M. Paiva, P. P. G. W. Rodrigues, M. I. P. Ferreira, and J. Lugon, “Disposal of waste from cementing operation from offshore oil and gas wells building,” *Ciência e Natura*, vol. 39, pp. 413–422, 5 2017.
- [7] K. Ogra, Y. Chandra, A. Pandey, V. Verma, A. Kumar, and R. K. Sinha, “Evolution of production logging technologies and capabilities: In quest to know the unknown, a brownfield case study, mumbai high, ongc western offshore,” in *SPE Oil and Gas India Conference and Exhibition*, SPE, 3 2012.
- [8] K. Manikonda, A. R. Hasan, N. H. Rahmani, O. Kaldirim, C. E. Obi, J. J. Schubert, and M. A. Rahman, “A gas kick model that uses the thermodynamic approach to account for gas solubility in synthetic-based mud,” in *SPE/IADC Middle East drilling Technology conference and exhibition*, p. D031S016R001, SPE, 2021.
- [9] “Peer review of the hot dry rock project at fenton hill, new mexico,” 12 1998.
- [10] K. C. Manger, “Assessing geothermal energy potential in upstate new york. final report, tasks 1, 3, and 4,” 7 1996.
- [11] S. Vij, S. Narasaiah, A. Walia, and G. Singh, “Multilaterals: An overview and issues involved in adopting this technology,” in *SPE India Oil and Gas Conference and Exhibition*, SPE, 2 1998.
- [12] J. W. Tester, G. Morris, R. Cummings, and R. Bivins, “Electricity from hot dry rock geothermal energy: technical and economic issues,” 1 1979.
- [13] P. Lienau and H. Ross, “Low-temperature resource assessment program. final report,” 2 1996.
- [14] Z. Hao, H. Chen, X. Jin, and Z. Liu, “Comparative study on the behavior of keyhole in analogy welding and real deep penetration laser welding,” *Materials (Basel, Switzerland)*, vol. 15, pp. 9001–9001, 12 2022.
- [15] C. E. Obi, Y. Falola, K. Manikonda, A. R. Hasan, and M. A. Rahman, “Flow pattern, pressure gradient relationship of gas kick under dynamic conditions,” in *Offshore Technology Conference*, p. D041S053R006, OTC, 2022.

- [16] R. L. Spross, “Halliburton sperry-sun doe high temperature lwd project,” 3 2005.
- [17] P. Laney, “Federal geothermal research program update - fiscal year 2004,” 3 2005.
- [18] J. Yang, J. Sun, Y. Bai, K. Lv, G. Zhang, and Y. Li, “Status and prospect of drilling fluid loss and lost circulation control technology in fractured formation.,” *Gels (Basel, Switzerland)*, vol. 8, pp. 260–260, 4 2022.
- [19] F. Benedict, R. Criss, M. Davisson, G. Eaton, G. Hudson, J. Kenneally, T. Rose, and D. Smith, “Hydrologic resources management program, fy 1998 progress report,” 7 1999.
- [20] M. Mendez, S. Garcia, R. Ahmed, H. Karami, M. Nasser, and I. Hussein, “Effect of fluid rheology on the performance of fibrous fluid in horizontal well cleanout,” in *SPE/ICoTA Well Intervention Conference and Exhibition*, SPE, 3 2022.
- [21] “Potential of hot-dry-rock geothermal energy in the eastern united states,” 11 1993.
- [22] D. Janiga, J. Kwaśnik, and P. Wojnarowski, “Utilization of discrete fracture network (dfn) in modelling and simulation of a horizontal well-doublet enhanced geothermal system (egs) with sensitivity analysis of key production parameters,” *Energies*, vol. 15, pp. 9020–9020, 11 2022.
- [23] “Eia model documentation: Documentation of the oil and gas supply module (ogsm),” 1 1997.
- [24] G. L. Morrison, R. Kroupa, A. Patil, J. Xu, S. Scott, and S. Olsen, “Experimental investigation of wellhead twin-screw pump for gas-well deliquification,” *Oil and Gas Facilities*, vol. 3, pp. 73–80, 3 2014.
- [25] J. C. Almeida, M. Rangel, G. Rey, M. Pacione, C. López, and A. Valdivieso, “Well of the future: A new approach to real-time wellbore stability monitoring in the colombian foothills,” *SPE/IADC Drilling Conference and Exhibition*, 3 2015.
- [26] A. P. Gilkey, C. W. Hansen, J. F. Schatz, D. K. Rudeen, and D. L. Lord, “Drspall :spallings model for the waste isolation pilot plant 2004 recertification,” 2 2006.
- [27] J. He and M. Luo, “Drilling rig hoisting platform security monitoring system design and application,” *Machines*, vol. 5, pp. 19–, 8 2017.
- [28] J. Sinner, “Improved miscible nitrogen flood performance utilizing advanced reservoir characterization and horizontal laterals in a class i reservoir - east binger (marchand) unit,” 6 2004.
- [29] K. Chhantyal, H. Viumdal, and S. Mylvaganam, “Soft sensing of non-newtonian fluid flow in open venturi channel using an array of ultrasonic level sensors—ai models and their validations,” *Sensors (Basel, Switzerland)*, vol. 17, pp. 2458–, 10 2017.
- [30] A. Hooshmandkoochi, M. Zaferanieh, and A. Malekzadeh, “Optimum technique selection for underbalanced drilling in iranian oil fields—a review of three major oil fields,” in *IADC/SPE Managed Pressure Drilling & Underbalanced Operations*, SPE, 3 2007.
- [31] J. T. Finger and B. J. Livesay, “Alternative wellbore lining methods: Problems and possibilities,” 8 2002.
- [32] G. Subbuswamy and X. Li, “Numerical study of film cooling enhancement in gas turbine combustor liner,” in *Volume 10: Heat Transfer, Fluid Flows, and Thermal Systems, Parts A, B, and C*, pp. 1237–1247, ASMEDC, 1 2008.
- [33] A. Mahmood, L. Ring, A. Alajbegovic, and S. Nienhuis, “Enabling new techniques: Nitrogen gas qualification and testing of a rotating control device,” in *IADC/SPE Managed Pressure Drilling & Underbalanced Operations Conference & Exhibition*, SPE, 3 2017.
- [34] C. E. Obi, Y. Falola, K. Manikonda, A. R. Hasan, and M. A. Rahman, “A machine learning analysis to relate flow pattern and pressure gradient during gas kicks under static conditions,” in *SPE Western Regional Meeting*, p. D011S006R002, SPE, 2022.
- [35] A. Wspanialy and M. Kyaw, “Surface drilling parameters and drilling optimization techniques: Are they useful tools in gas hydrate detection?,” *Energies*, vol. 15, pp. 4635–4635, 6 2022.
- [36] M. Horner, “Development of high-temperature turbine subsystem technology to a technology readiness status, phase ii. quarterly report, january-march 1981,” 4 1981.
- [37] J. Jiang, F. Rabbi, and Q. Wang, “Underbalanced drilling technology with air injection connector,” *Journal of Petroleum Exploration and Production Technology*, vol. 4, pp. 275–280, 11 2013.

- [38] J. B. Lewis, "The use of the computer and other special tools for monitoring a gas well blowout during the kill operation - offshore louisiana," *SPE Annual Fall Technical Conference and Exhibition*, 10 1977.
- [39] Z. Sun, X. Li, K. Wang, F. Wang, D. Chen, and Z. Li, "Determination of key technical parameters in the study of new pressure sealing technology for coal seam gas extraction.," *International journal of environmental research and public health*, vol. 19, pp. 4968–4968, 4 2022.
- [40] D. P. Howlett, C. Stewart, D. Reed, G. Begg, and C. A. Black, "Evolutionary well construction method that challenges convention," *SPE/IADC Drilling Conference*, 2 2007.
- [41] L. Peng, W. Kai, L. Bo, J. Yifeng, and G. Jianqiang, "Research on the effective influence radius of hydraulic reaming in mining seam," *The Open Fuels & Energy Science Journal*, vol. 8, pp. 161–167, 7 2015.