Original Research

Uncertainty-Aware Surface Pressure Inversion for Real-Time Kick Diagnosis in Managed-Pressure Drilling

Kasun Weerasinghe[1] and Tharindu Senavirathna[2]

[1]Lanka Institute of Information Sciences, Pelawatta Road, Anuradhapura 50000, Sri Lanka.
[2]Southern School of Computing, Weliwatta Road, Matara 81000, Sri Lanka.

Abstract
Reliable kick diagnosis remains challenging in modern drilling because the primary signals available at the surface are indirect, delayed, and entangled with operational transients. In managed-pressure drilling, the same actuation that improves safety by regulating annular pressure also complicates interpretation of pressure histories, since choke control, compressibility, and multiphase slip jointly shape the measured response. This paper proposes a technical framework that treats kick diagnosis as a constrained inverse problem in which surface pressure time series are used to infer a latent, depth-distributed multiphase state with explicit thermodynamic consistency and quantified uncertainty. The core contribution is a dual-timescale reduced-order model that couples a one-dimensional annular multiphase flow description with an equilibrium-constrained dissolved-gas component, while enforcing physically admissible mixture densities, gas fractions, and pressure gradients through projection operators. The reduced model is embedded in a sequential Bayesian inference loop that performs joint state estimation and online calibration of nuisance parameters representing frictional drift, sensor bias, and uncertain mixture rheology. A likelihood design is introduced that separates actuation-induced pressure changes from influx-induced signatures via a disturbance-aware innovation process, enabling probabilistic alarms that are robust to routine choke and pump adjustments. The framework also yields posterior distributions over influx size, type proxies, and migration speed, allowing risk-ranked control recommendations. Numerical experiments illustrate conditions under which the inversion is identifiable from surface pressure alone and quantify the computational budget required for real-time deployment.

1. Introduction

During drilling, the downhole annulus behaves as a distributed dynamical system whose states include pressure, phase fractions, velocities, temperature, and, when relevant, dissolved gas content in the liquid phase. Yet the measurement set is typically sparse, dominated by surface pressure, flow-out, standpipe pressure, and pit-volume indicators, each of which is influenced by equipment dynamics and operational decisions. Even when additional downhole sensors are present, they may be intermittent, delayed, or limited in spatial coverage. As a result, the most actionable and continuously available signals remain those at the surface, especially wellhead pressure and flow measurements. The difficulty is that these signals conflate multiple mechanisms: frictional losses, hydrostatic effects, annular compressibility, transient acceleration, and control actions such as choke modulation. A kick adds further complexity by introducing multiphase flow, slip, and evolving compressibility as gas expands during migration [1]. The central technical tension is therefore between the desire for fast, reliable alarms and the need for mechanistic interpretability that avoids brittle pattern matching.

Operationally, drillers and automated control systems must decide when a deviation in surface pressure is attributable to an influx rather than to benign transients such as pump ramps, cuttings loading, or changes in mud properties. The decision must be made quickly because time-to-surface decreases as gas expands and accelerates, and because well-control actions themselves change the observed signals. Traditional approaches rely on thresholds on pit gain, flow-out, or pressure trends, often supplemented

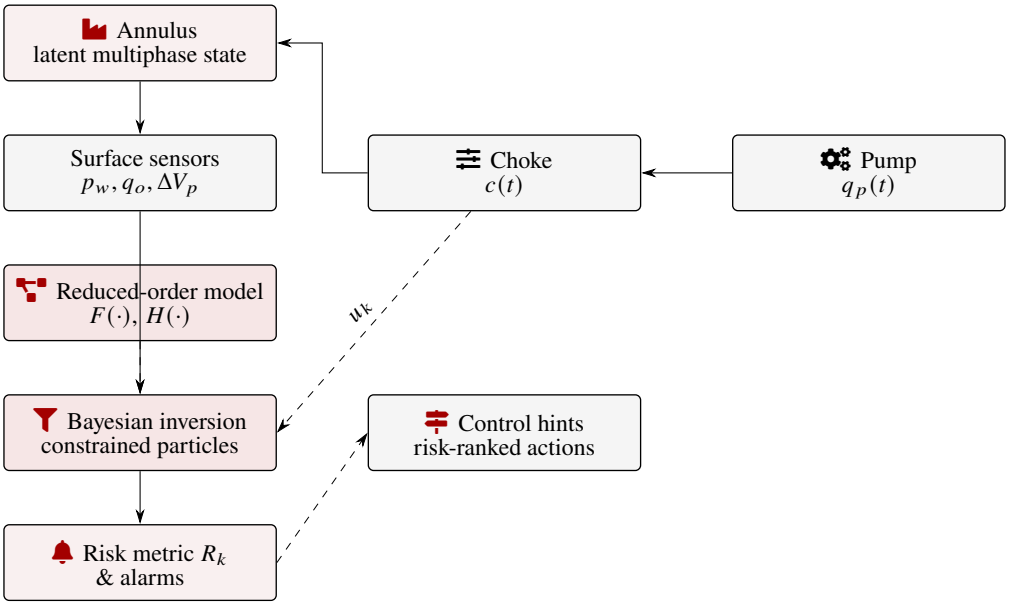

Figure 1: Closed-loop view of managed-pressure drilling diagnosis: surface measurements and known actuation are fused with a reduced-order multiphase model to produce posterior risk and decision outputs.

by heuristic rules. Such rules can be conservative, triggering nuisance alarms, or permissive, delaying detection. Data-driven approaches have demonstrated that surface pressure and flow signals contain discriminative structure for flow-regime identification and kick prognosis under controlled conditions, including during circulation where multiple supervised and unsupervised classifiers have achieved high accuracy in mapping surface measurements to annular flow patterns [2]. However, a key deployment concern is that purely discriminative models may underperform when exposed to conditions outside their training distribution, including different well geometries, mud formulations, temperature profiles, and control policies. This motivates hybrid strategies in which data-driven components are constrained by physics and deployed with calibrated uncertainty estimates.

A second deployment challenge is computational. High-fidelity multiphase simulators with detailed thermodynamics and spatial discretization can be accurate but are typically too slow for real-time inference when embedded in an estimator that must evaluate many candidate states per second. Conversely, simplified models can be fast but may omit critical couplings, such as the effect of gas solubility in synthetic-based muds, which alters effective compressibility, pit gain trajectories, and the pressure response during migration [3]. The appropriate modeling choice depends on the inference objective: if the goal is to estimate influx mass, depth, and evolution from pressure alone, the model must preserve the mapping from latent state to measured pressure with sufficient fidelity while remaining computationally tractable for sequential inversion.

This paper introduces a technical framework that addresses both challenges by combining a thermodynamically consistent reduced-order multiphase model with sequential Bayesian inversion specialized for surface pressure data. The paper's independent thesis is that reliable real-time kick diagnosis can be achieved by treating surface pressure as a noisy, control-contaminated observation of a constrained latent state, and by explicitly representing the uncertainty induced by unmodeled frictional drift, rheological variability, and sensor bias. The contribution is not a new classifier nor a reimplementation of existing kick simulators, but rather a model-inference architecture that converts mechanistic structure into probabilistic alarms and interpretable posterior estimates while meeting real-time computational budgets. The approach is designed to integrate naturally with managed-pressure drilling control loops

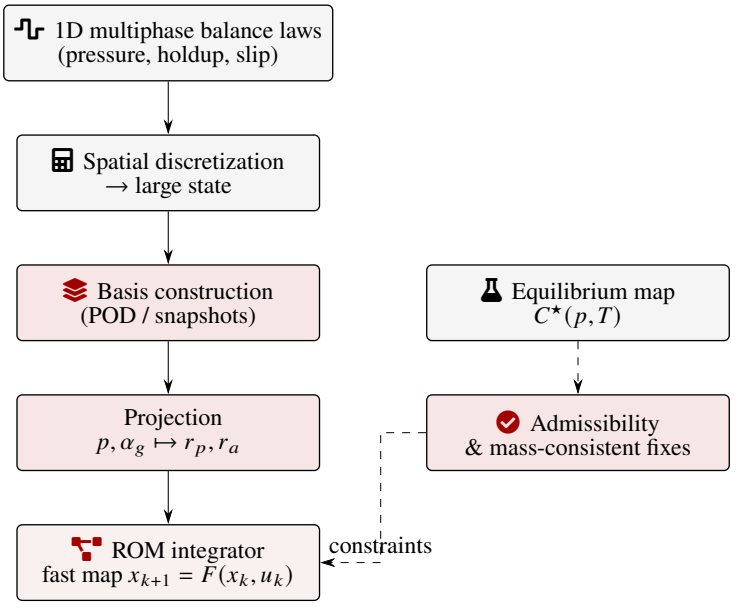

Figure 2: Reduced-order model pipeline: a projected representation retains the dominant pressure/holdup modes while enforcing thermodynamic admissibility through equilibrium-aware corrections.

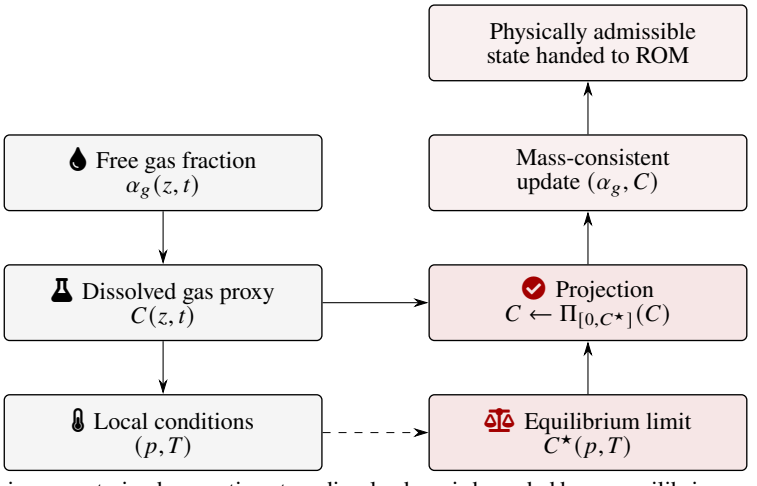

Figure 3: Equilibrium-constrained correction step: dissolved gas is bounded by an equilibrium map and reconciled with free-gas holdup through a mass-consistent projection.

by accounting for actuation in the inference model rather than treating it as a nuisance to be filtered out post hoc.

The remainder of the paper develops a consistent mathematical formulation and then presents the algorithmic machinery needed to deploy it [4]. The first technical step is to cast annular dynamics as a constrained state-space system with a measurement function that reflects the surface sensing configuration. The second step is to build a reduced-order model that retains the dominant pressure-flow couplings across the depth domain while respecting phase equilibrium constraints for dissolved gas. The third step is to perform sequential inference with a disturbance-aware likelihood that decouples controllable pressure changes from influx-driven anomalies. The final step is to define decision logic and validation protocols that translate posterior distributions into operationally meaningful alarms and estimates.

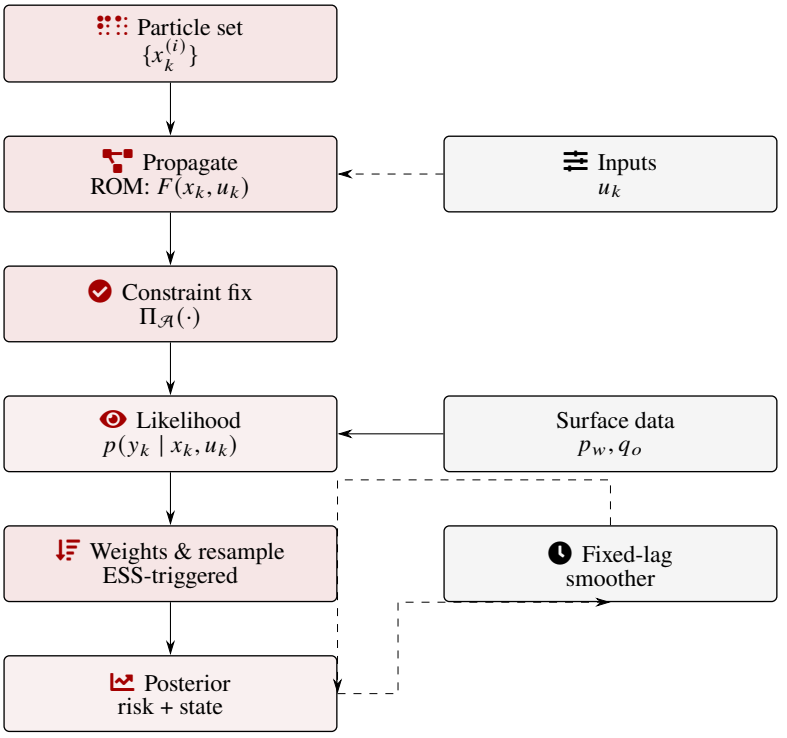

Figure 4: Constrained sequential inference: particles are propagated by the reduced model, projected to admissible multiphase states, and reweighted using a control-aware likelihood with optional short-lag smoothing.

Throughout, the design emphasizes identifiability, robustness to routine transients, and computational feasibility.

2. Problem Setting and Observability Structure

The wellbore annulus is modeled as a vertical or deviated conduit with slowly varying cross-sectional area, possibly including tool joints, eccentricity effects, and changes in hydraulic diameter [5]. For inference, the essential question is what aspects of the distributed state can be recovered from surface measurements. Let $z \in [0, L]$ denote measured depth along the annulus, with $z = 0$ at the wellhead and $z = L$ at the bit or the deepest modeled point. The latent state includes pressure $p(z, t)$, mixture velocity $u_m(z, t)$, gas volume fraction $\alpha_g(z, t)$, and additional internal variables required to represent dissolved gas and mixture properties. The observation set typically includes wellhead pressure $p_w(t) = p(0, t)$, choke position $c(t)$, pump rate command $q_p(t)$, measured flow-out $q_o(t)$, and pit-volume changes $\Delta V_p(t)$ when available. In managed-pressure drilling, $p_w(t)$ is strongly influenced by choke-induced backpressure and by the compressibility of the annulus-fluid system. The same pressure signal contains the signature of an influx, but the signature is not unique without additional structure [6].

To formalize this, consider a generic nonlinear state-space model,

$$\dot{x}(t) = f(x(t), u(t), \theta) + w(t), \quad x(0) = x_0, \tag{2.1}$$

$$y(t) = h(x(t), u(t), \theta) + v(t), \tag{2.2}$$

where $x(t)$ represents a finite-dimensional parameterization of the distributed annulus state, $u(t)$ denotes known inputs such as pump rate and choke command, θ collects uncertain parameters, $w(t)$ is process noise capturing model mismatch, and $v(t)$ is measurement noise. The measurement $y(t)$ may include

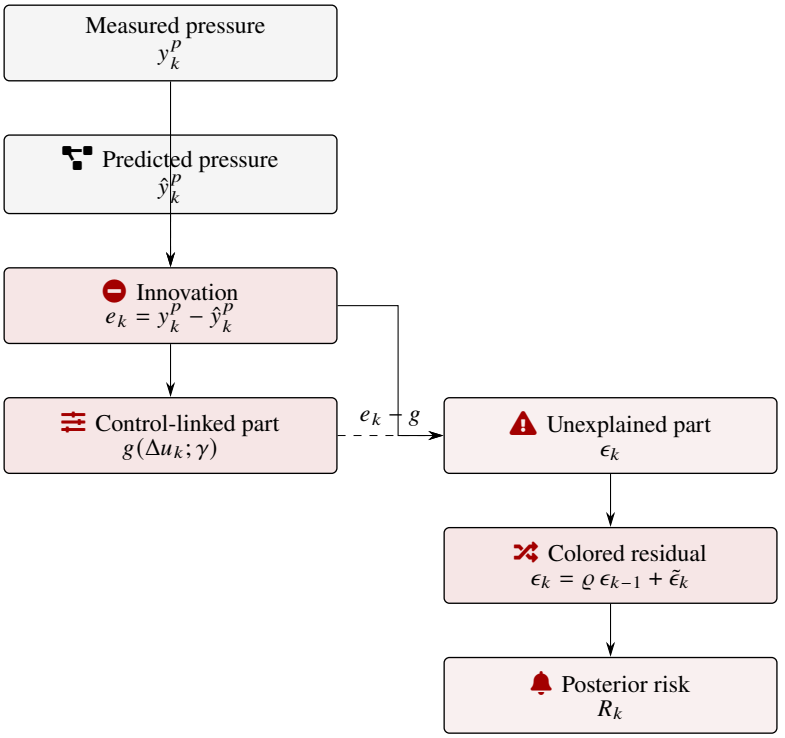

Figure 5: Disturbance-aware innovation structure: control-correlated residual components are separated from unexplained deviations, enabling conservative yet sensitive probabilistic kick risk updates.

multiple channels, but the design here focuses on cases where $p_w(t)$ is the primary high-rate observation and other channels are lower-rate or missing.

The key difficulty is that the mapping from $x(t)$ to $p_w(t)$ is many-to-one if $x(t)$ is too rich. Therefore, a principled reduction must preserve the subspace of state variations that are observable from surface pressure under realistic actuation [7]. Observability is influenced by both the physics and the input design. Intuitively, pressure changes induced by choke modulation propagate down the annulus and interact with compressibility; these controlled perturbations can serve as probing signals that improve identifiability if modeled correctly. However, they also introduce confounding if treated as exogenous disturbances. The proposed framework treats choke and pump commands as known inputs and explicitly models the resulting pressure response so that residual deviations can be attributed probabilistically to influx dynamics.

Kick phenomena introduce additional latent variables. At a minimum, one may define an influx source term at depth z_k with an onset time t_k, characterized by an equivalent mass influx rate $\dot{m}_k(t)$ over a duration τ_k. The influx type influences compressibility, dissolution, and slip, but in field scenarios the exact composition is rarely known [8]. Instead, the inference target is a set of proxies that capture how the influx behaves in the annulus: effective gas density, solubility tendency, and expansion law parameters. The inference model must accommodate these proxies without exploding the state dimension.

A crucial modeling choice concerns the representation of phase behavior. In water-based muds, dissolution of gas may be limited, but in synthetic-based muds, solubility can be significant and can delay the appearance of free gas while still changing the effective compressibility and volume balance. Thermodynamic treatment of solubility has been shown to materially affect predicted dissolved-gas fraction and bubble-point location in mechanistic kick modeling for synthetic-based muds, with a thermodynamic equilibrium approach yielding close agreement with commercial process simulation for key solubility quantities [9]. For inference, this implies that a surface-pressure-only estimator that ignores dissolution

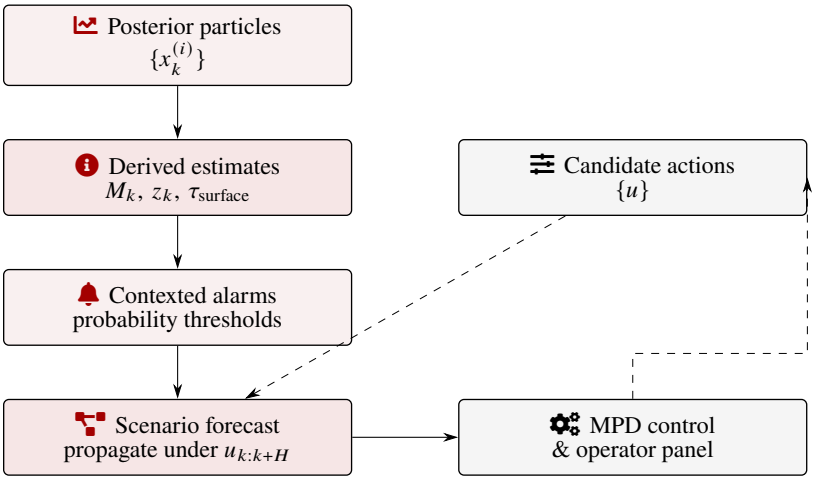

Figure 6: Operational outputs: the posterior supports uncertainty-bounded influx descriptors, probabilistic alarms, and fast forward forecasts under candidate choke/pump actions for control-compatible decision-making.

may systematically bias estimated influx size and migration speed, because part of the influx mass may initially be stored in the liquid phase rather than appearing as a free gas fraction.

Observability is also affected by flow regime. The relation between pressure gradient and flow pattern can change sharply when transitioning among dispersed bubbles [10], slug flow, churn, and annular regimes. Under shut-in conditions [11] [12], the system is nearly static except for gas migration and expansion, whereas under circulation the liquid momentum dominates but can entrain and redistribute gas. The pressure signal may contain regime-dependent features such as oscillations, slope changes, and transient damping. Experiments under static conditions have indicated that surface pressure responses and gas evolution data can be used by machine learning models to detect kick symptoms with high accuracy, and that inferred pressure-gradient patterns correlate with observed flow regimes and influx size [13]. While the present paper does not aim to replicate those pattern-recognition results, it leverages the underlying insight: the pressure signal carries regime information, and therefore a reduced-order model must allow for regime-induced changes in effective friction and slip that can be inferred online.

Given these considerations, the problem is posed as follows. Given input histories $u(t)$ and surface measurements $y(t)$ dominated by $p_w(t)$, estimate in real time the posterior distribution of (i) latent annulus state variables that determine pressure propagation and gas migration, (ii) influx parameters such as onset time, duration, and effective mass rate, and (iii) nuisance parameters representing frictional drift, sensor bias, and rheology variability. The estimator must run at a cadence compatible with field data rates, often 1–10 Hz, and must produce alarms with controlled false-positive rates while providing interpretable uncertainty bounds.

To make this problem tractable, the architecture adopts three structural assumptions that are technical rather than heuristic [14]. First, the distributed state is projected onto a low-dimensional manifold that preserves the dominant modes of pressure and gas-fraction variation observable at the surface. Second, phase equilibrium for dissolved gas is enforced as a constraint rather than approximated by an unconstrained surrogate, so that inferred states remain physically admissible. Third, nuisance parameters are treated as slowly varying latent variables with explicit priors, allowing the estimator to adapt to drift without mistaking it for an influx. These assumptions guide the model construction in the next section.

3. Thermodynamically Consistent Multiphase Reduced-Order Model

A mechanistic model for annular kick dynamics must represent compressibility, slip between phases, and pressure losses, while remaining computationally manageable. The starting point is a one-dimensional

Variable	Symbol	Units	Notes
Pressure	$p(z,t)$	Pa	Annular pressure
Mixture velocity	$u_m(z,t)$	m/s	Depth dependent
Gas fraction	$\alpha_g(z,t)$	–	Free gas holdup
Liquid fraction	$\alpha_\ell(z,t)$	–	$1-\alpha_g$
Dissolved gas	$C(z,t)$	mol/kg	In liquid phase
Mixture density	ρ_m	kg/m^3	α-weighted

Table 1: State variables summary.

Channel	Symbol	Type	Rate
Wellhead pressure	$p_w(t)$	Meas.	1–10 Hz
Flow-out	$q_o(t)$	Meas.	< 2 Hz
Pit volume	$\Delta V_p(t)$	Meas.	Slow
Pump rate	$q_p(t)$	Input	1–10 Hz
Choke opening	$c(t)$	Input	1–10 Hz

Table 2: Surface channels overview.

Balance	Representative form	Key terms	Kick impact
Gas mass	$\partial_t(\alpha_g\rho_g) + \partial_z(\alpha_g\rho_g u_g)$	Slip, source S_k	Influx entry
Liquid mass	$\partial_t(\alpha_\ell\rho_\ell) + \partial_z(\alpha_\ell\rho_\ell u_\ell)$	Γ_d	Solubility
Dissolved gas	$\partial_t(\alpha_\ell\rho_\ell C)$	$\Gamma_d - \Gamma_e$	Bubble point
Momentum	$\partial_t(\rho_m u_m) + \partial_z(\rho_m u_m^2)$	p_z, friction	Pressure waves

Table 3: Core conservation equations.

Block	Symbol	Dim.	Typical size
Pressure modes	r_p	r_p	3–6
Gas-fraction modes	r_a	r_a	3–6
Dissolved-gas modes	η	r_η	1–3
Nuisance vector	ξ	d_ξ	4–10

Table 4: Reduced-order blocks.

conservation framework for a gas-liquid mixture flowing in the annulus [15]. Let $\alpha_g(z,t)$ be the volumetric gas fraction of the free gas phase, and $\alpha_\ell(z,t) = 1-\alpha_g(z,t)$ the liquid fraction. Let $\rho_g(p,T)$ and $\rho_\ell(p,T,C)$ denote densities of gas and liquid, with the liquid density depending additionally on dissolved gas concentration $C(z,t)$ or an equivalent dissolved mole fraction. Let u_g and u_ℓ denote phase velocities, and define mixture superficial velocity $u_m = \alpha_g u_g + \alpha_\ell u_\ell$. Temperature $T(z,t)$ may be prescribed from a geothermal profile with slow dynamics relative to pressure transients, or it may be included with a simplified energy balance if needed; for surface-pressure inversion at short time scales, treating T as known is often sufficient, but the formulation permits uncertainty in T through θ.

Parameter	Symbol	Units	Note
Mass-rate bins	$m_{k,j}$	kg/s	Windowed influx
Onset time	t_k^{on}	s	Start of entry
Entry depth	z_k	m	Source location
Duration scale	τ_k	s	Effective length
Gas density proxy	$\rho_{g,k}$	kg/m^3	Compressibility

Table 5: Influx parameterization.

Group	Example	Symbol	Drift
Friction	Multiplier	κ_f	Random walk
Choke map	Coefficients	θ_c	Random walk
Slip law	β-vector	β_i	Random walk
Sensor bias	Offset	b_p	Mean-revert
Solubility	Relaxation	τ_{rel}	Random walk

Table 6: Nuisance parameters.

Term	Symbol	Definition	Role
Raw innovation	e_k	$y_k^p - \hat{y}_k^p$	Total residual
Control part	$g(\cdot)$	$g(u_k, u_{k-1}; \gamma)$	Actuation fit
Anomaly part	ϵ_k	$e_k - g(\cdot)$	Likelihood core
AR factor	ϱ	$\epsilon_k = \varrho\epsilon_{k-1} + \tilde{\epsilon}_k$	Colored noise
Noise level	σ_p^2	$\mathrm{Var}(\tilde{\epsilon}_k)$	Sensor spread

Table 7: Innovation decomposition.

Quantity	Symbol	Meaning	Use
Kick risk	R_k	$\mathbb{P}(\sum_j m_{k,j} > m_{\min})$	Alarm trigger
Total influx mass	M_k	Integrated influx	Severity level
Free-gas volume	$V_{g,k}$	Gas in annulus	Migration risk
Time-to-surface	τ	Arrival time	Action timing
Context threshold	$\pi(s_k)$	Risk limit	Mode-dependent

Table 8: Estimator output metrics.

A generic two-phase mass conservation model with dissolution is

$$\partial_t\big(\alpha_g\rho_g A\big) + \partial_z\big(\alpha_g\rho_g u_g A\big) = -\Gamma_d(p,T,C) + S_k(z,t), \tag{3.1}$$

$$\partial_t(\alpha_\ell\rho_\ell A) + \partial_z(\alpha_\ell\rho_\ell u_\ell A) = \Gamma_d(p,T,C), \tag{3.2}$$

where $A(z)$ is cross-sectional area, $S_k(z,t)$ represents an influx source localized near $z = z_k$ during a limited time window, and Γ_d is the dissolution mass transfer term from free gas to dissolved gas in the liquid, with sign convention such that $\Gamma_d > 0$ transfers mass to the liquid phase [16]. The dissolved gas

concentration evolves according to

$$\partial_t(\alpha_\ell \rho_\ell C A) + \partial_z(\alpha_\ell \rho_\ell C u_\ell A) = \Gamma_d(p,T,C) - \Gamma_e(p,T,C), \tag{3.3}$$

where Γ_e represents exsolution when the dissolved fraction exceeds equilibrium at local conditions. In many drilling contexts, dissolution and exsolution are fast relative to advective time scales, motivating a quasi-equilibrium constraint rather than explicit kinetics. The model here adopts a constrained-equilibrium approach: C is projected toward an equilibrium manifold $C^\star(p,T)$ while preserving mass balance. This replaces uncertain kinetic parameters with a relaxation time scale that can be treated as a nuisance parameter.

Momentum balance can be represented in several ways, depending on fidelity. For computational efficiency, the reduced-order model does not resolve detailed wave propagation at acoustic speeds. Instead, it uses a low-Mach, compressible mixture momentum approximation that captures the dominant relationship between pressure gradient, mixture acceleration, and friction [17]. A representative form is

$$\partial_t(\rho_m u_m A) + \partial_z\left(\rho_m u_m^2 A\right) + A\,\partial_z p = -A\,\rho_m g \sin\phi(z) - A\,\mathcal{F}(u_m, \alpha_g, \mu_{\text{eff}}, D_h) + \mathcal{U}(z,t), \tag{3.4}$$

where $\rho_m = \alpha_g \rho_g + \alpha_\ell \rho_\ell$ is mixture density, $\phi(z)$ is deviation angle, $\mathcal{F}$ is a friction term dependent on effective viscosity μ_{eff} and hydraulic diameter $D_h(z)$, and $\mathcal{U}$ accounts for distributed forcing induced by pump and choke actions through boundary conditions. Slip is introduced through a closure relating $u_g - u_\ell$ to local conditions. Instead of selecting a single regime-specific correlation, the framework uses a parametric slip law with online adaptation:

$$u_g - u_\ell = \beta_0 + \beta_1\, u_m + \beta_2 \sqrt{\frac{g D_h |\rho_\ell - \rho_g|}{\rho_m}}\, \Psi(\alpha_g), \tag{3.5}$$

where $\Psi(\alpha_g)$ shapes the transition from dispersed to slug-like behavior, and β_i are nuisance parameters that drift slowly and are inferred sequentially. This is not intended as a universal slip model; rather, it provides a flexible representation that can capture regime shifts as changes in effective parameters without requiring explicit regime classification.

The equilibrium constraint for dissolved gas is expressed as an algebraic relation between C and (p,T) [18]. For inference, it is convenient to represent the constraint through an equilibrium operator $\mathcal{E}$:

$$C^\star(p,T) = \mathcal{E}(p,T;\theta_s), \tag{3.6}$$

where θ_s includes uncertain solubility parameters, potentially including an effective partition coefficient that can depend on mud formulation. The reduced-order model enforces

$$C \leftarrow \Pi_{\mathcal{M}}(C), \quad \mathcal{M} = \left\{C : 0 \le C \le C^\star(p,T)\right\}, \tag{3.7}$$

where $\Pi_{\mathcal{M}}$ is a projection that ensures the inferred dissolved fraction does not exceed equilibrium for given (p,T), and that it remains nonnegative. When mass balance would be violated by direct projection, the correction is accompanied by a compensating adjustment to the free gas fraction α_g such that total gas mass is conserved locally. This yields a thermodynamically consistent correction step that is stable under sequential inference.

The full distributed model is still too large for real-time inversion if discretized finely. The core reduction step is to identify a low-dimensional basis for the dominant variations in pressure and gas fraction that influence the wellhead pressure [19]. Let $p(z,t)$ and $\alpha_g(z,t)$ be discretized over N depth nodes, forming vectors $p_N(t) \in \mathbb{R}^N$ and $a_N(t) \in \mathbb{R}^N$. Define centered fields $\tilde{p}_N(t) = p_N(t) - \bar{p}_N$

and $\tilde{a}_N(t) = a_N(t) - \bar{a}_N$, where bars denote nominal profiles under normal drilling. Construct basis matrices $\Phi_p \in \mathbb{R}^{N\times r_p}$ and $\Phi_a \in \mathbb{R}^{N\times r_a}$ with $r_p, r_a \ll N$, such that

$$\tilde{p}_N(t) \approx \Phi_p r_p(t), \tag{3.8}$$

$$\tilde{a}_N(t) \approx \Phi_a r_a(t), \tag{3.9}$$

where $r_p(t)$ and $r_a(t)$ are reduced coordinates. The basis may be obtained from offline simulations spanning plausible operating ranges, from lab data, or from online incremental updates. The key requirement is that the basis includes modes that affect $p(0,t)$ strongly, including long-wavelength pressure shifts and localized perturbations that can propagate upward. Because the primary measurement is at the boundary $z = 0$, modes that are nearly orthogonal to the wellhead observation are not useful and can be omitted.

To couple pressure and gas fraction dynamics, define a combined reduced state

$$x(t) = \begin{bmatrix} r_p(t) \\ r_a(t) \\ \eta(t) \\ \xi(t) \end{bmatrix}, \tag{3.10}$$

where $\eta(t)$ collects reduced variables for dissolved gas and potentially mixture compressibility, and $\xi(t)$ contains nuisance parameters such as friction and slip coefficients and sensor bias [20]. The reduced dynamics are obtained by a Petrov–Galerkin projection of the discretized conservation equations onto the chosen bases, resulting in a nonlinear ordinary differential equation system:

$$\dot{x}(t) = f_r(x(t), u(t)) + w_r(t). \tag{3.11}$$

In practice, direct Galerkin projection of strongly nonlinear two-phase closures can lead to instability. The framework therefore uses structure-preserving reduction: advection and source terms are evaluated in a lifted space using sparse sampling operators, and the result is projected back to the reduced coordinates. The projection is accompanied by constraint enforcement to maintain admissibility, including $0 \le \alpha_g \le 1$, bounded dissolved fraction, and nonnegative densities.

Boundary conditions reflect pump inflow at the bottom and choke-controlled outflow at the top. At the wellhead, the choke introduces a relation between flow-out and wellhead pressure [21]. Rather than modeling choke hydraulics in full detail, which may depend on equipment-specific characteristics, a parametric choke map is introduced:

$$q_o(t) = C(p_w(t), c(t); \theta_c)\,, \tag{3.12}$$

where θ_c captures uncertain choke coefficients that can be calibrated online. This is important because choke behavior is a major confounder in surface pressure analysis; if the choke model is wrong, the estimator may attribute systematic discrepancies to influx dynamics. The calibration of θ_c is treated as part of $\xi(t)$ with a slow drift model.

The reduced-order model thus produces predicted wellhead pressure $\hat{p}_w(t)$ as a function of reduced state and inputs, and it produces predictions for auxiliary channels such as flow-out when available. Most importantly, it provides a computationally cheap map from candidate influx scenarios to surface pressure trajectories, enabling the sequential inversion described next.

4. Sequential Bayesian Inference with Disturbance-Aware Likelihood

The inference engine must ingest streaming measurements and update beliefs about the latent state and influx parameters [22]. A central design requirement is robustness to routine operations. Pump ramps, choke adjustments, and changes in drilling conditions can produce pressure transients that resemble

influx signatures. The proposed estimator therefore constructs a likelihood that explicitly accounts for control actions and for systematic model errors that correlate with those actions.

Let the discrete-time sampling interval be Δt, with times $t_k = k\Delta t$. Define the reduced state $x_k = x(t_k)$ and inputs $u_k = u(t_k)$, and consider the stochastic evolution

$$x_{k+1} = F(x_k, u_k) + \omega_k, \tag{4.1}$$

$$y_k = H(x_k, u_k) + \nu_k, \text{ [23]} \tag{4.2}$$

where F and H are discrete-time maps obtained from integrating the reduced-order model over Δt, and ω_k, ν_k are process and measurement noise. The observation y_k includes at least $p_w(t_k)$ and optionally $q_o(t_k)$ and pit-volume change. The unknown influx enters through an augmented parameter vector $\theta_k^{(k)}$ that represents onset, depth, and mass rate. Rather than attempting to estimate a continuous influx-rate function directly, the estimator uses a parsimonious parameterization. One convenient choice is a piecewise-constant influx rate over a sliding window:

$$\dot{m}_k(t) \approx \sum_{j=0}^{J-1} m_{k,j}\, \mathbf{1}_{[t_{k-j-1}, t_{k-j})}(t), \tag{4.3}$$

where $m_{k,j}$ are nonnegative coefficients representing recent influx contributions and are part of the augmented state with a sparsity-inducing prior. This allows the estimator to represent unknown onset times and durations without discrete change-point enumeration.

The posterior of interest is [24]

$$p(x_k, \theta_k \mid y_{1:k}, u_{1:k}), \tag{4.4}$$

where θ_k collects influx coefficients and nuisance parameters. Exact inference is intractable due to nonlinearity, constraints, and non-Gaussian noise. The framework therefore uses a constrained particle filter with Rao–Blackwellization: continuous approximately linear components are updated with Kalman-type recursions, while strongly nonlinear and constrained components are sampled. The computational benefit is that a modest number of particles can suffice when the reduced-order model has low dimension and when nuisance parameters are modeled with smooth priors.

A critical innovation is the disturbance-aware likelihood. Let y_k^p denote measured wellhead pressure and $\hat{y}_k^p$ the model prediction. A naive likelihood assumes innovations $e_k = y_k^p - \hat{y}_k^p$ are independent Gaussian noise. In practice, innovations are correlated with control actions, especially choke motion, because unmodeled choke dynamics and compressibility effects manifest as systematic residuals [25]. To mitigate this, define an innovation decomposition:

$$e_k = \underbrace{g(u_k, u_{k-1}; \gamma)}_{\text{control-correlated residual}} + \epsilon_k, \tag{4.5}$$

where g is a low-dimensional regression model of residuals on control increments, parameterized by γ, and ϵ_k is the residual component attributed to unmodeled physics including potential influx. The parameters γ are included in the nuisance state and adapted online. The likelihood is then built on ϵ_k rather than e_k:

$$p(y_k^p \mid x_k, u_k, \gamma) \propto \exp\left(-\frac{1}{2} \frac{\left(y_k^p - \hat{y}_k^p - g(u_k, u_{k-1}; \gamma) \right)^2}{\sigma_p^2} \right). \tag{4.6}$$

This construction is deliberately conservative: it allocates part of the discrepancy to control-correlated effects, reducing false alarms during routine adjustments. When an influx occurs, the residual will

exhibit a component that cannot be explained by control increments alone, pushing ϵ_k beyond its typical distribution and thereby shifting posterior mass toward nonzero influx coefficients [26].

To incorporate constraints, each particle state is projected onto an admissible set after propagation:

$$x_{k+1}^{(i)} \leftarrow \Pi_{\mathcal{A}}\left(x_{k+1}^{(i)}\right), \tag{4.7}$$

where $\mathcal{A}$ enforces bounds on gas fraction, dissolved fraction, densities, and pressure monotonicity with depth when appropriate. Projection is performed in the physical state space, then mapped back to reduced coordinates. This is important because unconstrained particles can produce nonphysical predictions that corrupt the likelihood and cause particle collapse.

Influx coefficients are assigned priors that encourage sparsity and nonnegativity. A convenient continuous prior is a rectified Laplace-like distribution implemented through a latent scale mixture:

$$m_{k,j} \mid \lambda_{k,j} \sim \text{HalfNormal}(\lambda_{k,j}), \quad \lambda_{k,j} \sim \text{Exponential}(\tau^{-1}), \tag{4.8}$$

where τ controls sparsity [27]. This prior favors $m_{k,j}$ near zero under normal drilling, but allows rapid growth when the likelihood supports an influx. Because the estimator operates online, τ can be tuned to achieve desired false-positive behavior, and it can be adapted based on recent innovation statistics.

The estimator also maintains a probabilistic alarm variable. Define a scalar risk metric R_k as the posterior probability that influx mass rate exceeds a threshold over the recent window:

$$R_k = \mathbb{P}\left(\sum_{j=0}^{J-1} m_{k,j} > m_{\min} \,\middle|\, y_{1:k}, u_{1:k}\right). \tag{4.9}$$

An alarm is triggered when R_k exceeds a user-specified confidence level. The advantage of this probabilistic design is that it provides graded information: a small but increasing R_k can prompt heightened monitoring without immediate shutdown, while a large R_k can trigger decisive well-control actions. The threshold can be selected based on acceptable risk and on the cost of false alarms, and it can be context-dependent, for example stricter during connections when flow is stopped [28].

Beyond alarm triggering, the estimator yields posterior distributions over derived quantities such as free gas volume in the annulus, expected migration speed, and predicted time to reach the wellhead. These are computed by propagating particles forward under candidate control policies. Importantly, the forward propagation uses the same reduced-order model and constraint projections, ensuring consistency between inference and prediction.

Computationally, the dominant cost is evaluating $F(x_k, u_k)$ for each particle. The reduced-order model is designed so that this evaluation is inexpensive. Additionally, the estimator uses an adaptive resampling schedule: resampling is performed only when effective sample size drops below a threshold [29]. To avoid sample impoverishment, a rejuvenation step adds small, constraint-respecting perturbations to nuisance parameters, and the influx coefficients are updated with a local Metropolis step that preserves nonnegativity. These details are essential for stable operation over long drilling intervals.

A further refinement is a smoothing step over a short lag, which improves onset-time estimation. Because surface pressure responds to an influx with delay, filtering alone can lag detection. A fixed-lag smoother updates the posterior of recent influx coefficients using future observations within a short horizon. The horizon is selected so that it remains operationally useful and does not delay alarms excessively. The smoother improves parameter identifiability by using the full shape of the transient rather than only its early portion [30].

5. Identifiability, Error Modes, and Physics-Guided Regularization

A surface-pressure-only inversion problem is only solvable when the mapping from latent influx parameters to pressure trajectories is sufficiently distinctive relative to nuisance variations. This section analyzes identifiability in a pragmatic sense: not whether the model is globally identifiable in an abstract mathematical sense, but whether the posterior can concentrate rapidly enough to support operational decisions.

Consider a simplified linearization around a nominal trajectory without influx. Let δx_k denote deviations in reduced state and δm_k deviations in influx coefficients. Linearizing the reduced dynamics and measurement yields

$$\delta x_{k+1} = A_k\,\delta x_k + B_k\,\delta u_k + G_k\,\delta m_k + \omega_k, \tag{5.1}$$

$$\delta y_k = C_k\,\delta x_k + D_k\,\delta u_k + \nu_k, \text{ [31]} \tag{5.2}$$

where matrices are Jacobians evaluated along the nominal path. The question is whether δm_k can be inferred given $\delta y_{1:k}$ and known $\delta u_{1:k}$. A necessary condition is that the impulse response from δm to δy is not in the span of responses from nuisance parameter variations and control-correlated residuals. In practice, this condition fails when friction drift or choke-map error can mimic the pressure trajectory induced by a small influx. The framework addresses this with two regularization strategies rooted in physics.

First, nuisance parameters are modeled as slow processes. For example, friction multiplier κ_f and choke coefficient κ_c are assigned random-walk dynamics with small variance: [32]

$$\kappa_{f,k+1} = \kappa_{f,k} + \zeta_{f,k}, \quad \zeta_{f,k} \sim \mathcal{N}(0, \sigma_f^2), \tag{5.3}$$

$$\kappa_{c,k+1} = \kappa_{c,k} + \zeta_{c,k}, \quad \zeta_{c,k} \sim \mathcal{N}(0, \sigma_c^2). \tag{5.4}$$

This prevents the estimator from explaining fast pressure deviations by rapid changes in nuisance parameters, thereby reserving fast residual structure for influx-related hypotheses. The variances σ_f^2, σ_c^2 encode operational belief: mud rheology and choke behavior do drift, but not abruptly at high frequency under steady operations. If operations include known events that can change these parameters rapidly, such as large changes in mud weight or equipment switching, the model can temporarily increase these variances based on event flags.

Second, the thermodynamic constraint provides a regularization on how gas mass can manifest in the pressure response. When gas can dissolve, early pressure changes may reflect effective compressibility increases without immediate growth in free gas fraction. An estimator that ignores dissolution might infer a free gas fraction that is inconsistent with pressure behavior and pit gain [33]. By enforcing an equilibrium-constrained dissolved component, the estimator narrows the set of admissible explanations for a given pressure history. This is not a claim that equilibrium holds exactly, but rather that it is a plausible fast manifold relative to the measurement cadence, and deviations from it can be absorbed into the relaxation parameter included in the nuisance state.

Error modes are analyzed through scenario-based reasoning. One common error mode is a false alarm triggered by pump transients. Pump rate changes alter annular friction and can induce compressibility-related pressure changes even without influx. The disturbance-aware likelihood explicitly accounts for control increments, but if the pump and choke dynamics are poorly modeled, residuals may persist [34]. The framework mitigates this by treating model mismatch as colored noise through an autoregressive innovation model:

$$\epsilon_k = \varrho\,\epsilon_{k-1} + \tilde{\epsilon}_k, \quad \tilde{\epsilon}_k \sim \mathcal{N}(0, \sigma_\epsilon^2), \tag{5.5}$$

with ϱ inferred online. This allows the estimator to represent slowly decaying residuals from unmodeled dynamics without incorrectly attributing them to influx. A second error mode is missed detection during connections or shut-in periods when flow is stopped. In such cases, pressure dynamics may be dominated

by gas migration and expansion, and the signal may be subtle. The estimator addresses this by switching model regimes: when pump rate is near zero, the reduced model emphasizes buoyant migration and compressibility, while the nuisance friction parameters are down-weighted because frictional losses are minimal [35]. This regime switching is input-driven and does not require explicit classification; it emerges from the dependence of the reduced dynamics on u_k.

A third error mode concerns geometric uncertainty, such as unknown washouts or eccentricity changes that alter hydraulic diameter and hence friction and holdup. The framework treats geometric uncertainty as part of the nuisance parameter vector with spatially aggregated effects. While this cannot recover detailed geometry, it can capture its first-order influence on pressure gradients. Importantly, the posterior uncertainty will reflect the ambiguity: if geometry uncertainty is large, the estimator will produce wider credible intervals for influx size and may delay a high-confidence alarm unless the pressure signature is strong.

Identifiability is improved when the control policy provides excitation [36]. In managed-pressure drilling, small, deliberate choke perturbations can probe annular compliance and reveal whether observed pressure changes are consistent with a purely liquid-filled annulus or with an evolving gas fraction. The proposed framework can exploit such probing because control actions are included in u_k and modeled explicitly. This creates a path toward active diagnosis: rather than passively waiting for an unambiguous signature, the control system can apply a small perturbation designed to maximize expected information gain about the influx parameters while respecting operational constraints. Formally, one can define an expected information gain objective based on the posterior predictive distribution, but the present paper focuses on the inference machinery required to support such objectives.

6. Real-Time Decision Outputs and Control-Compatible Forecasting

An estimator is operationally useful only if its outputs map cleanly to decisions. The proposed framework produces three classes of outputs: probabilistic alarms, parameter estimates with uncertainty, and forward forecasts under candidate control actions [37]. This section describes how these outputs are generated and how they can be consumed by a managed-pressure drilling control system.

The primary alarm output is the posterior risk metric R_k defined earlier. However, a single threshold on R_k may not be sufficient because operational contexts differ. During steady drilling, a moderate risk may justify a gentle increase in monitoring, whereas during tripping or connections, the tolerance for risk may be lower. The framework therefore defines context-conditioned alarm thresholds through a policy function π:

$$\text{Alarm}_k = \mathbf{1}\{R_k > \pi(s_k)\}\,, \tag{6.1}$$

where s_k represents operational state indicators such as pump status, connection mode, and choke authority [38]. The policy can be configured conservatively without modifying the estimator itself. Because R_k is a probability, the policy has a direct interpretation, unlike ad hoc thresholding on raw pressure derivatives.

The second output class includes posterior summaries of influx descriptors. Let M_k denote inferred total influx mass over the window, and let z_k denote inferred depth of entry if included in the model. The estimator produces $\mathbb{E}[M_k]$, credible intervals, and, when identifiable, multimodal distributions reflecting ambiguity among depths. Rather than forcing a single point estimate, the framework preserves distributional information that can guide cautious decision-making. For example, if depth is uncertain but influx mass is likely nonzero, a control action that increases bottomhole pressure across a broad depth range may be preferred [39].

The third output class is control-compatible forecasting. Given a candidate control sequence $u_{k:k+H}$ over a horizon H, the estimator propagates particles forward to generate a posterior predictive distribution for wellhead pressure and other quantities. This supports questions such as whether a planned choke

adjustment is likely to suppress further influx or whether it may accelerate gas migration by reducing hydrostatic pressure. Forecasting is also essential for estimating time-to-surface. Define a derived event time $\tau^{(i)}$ for each particle as the time when the particle's free gas fraction reaches the wellhead beyond a small threshold. The distribution of τ provides a probabilistic time-to-surface metric.

To support real-time use, forecasting must be fast. The reduced-order model enables propagation of hundreds of particles over a horizon of minutes within seconds on typical industrial computing hardware, assuming efficient implementation [12]. More importantly, the reduced-order model is stable under forward integration because constraint projection prevents unphysical blow-up. This stability is crucial when forecasts are used to rank control actions; unstable trajectories would yield misleading risk estimates.

Integration with control loops requires careful attention to feedback. In managed-pressure drilling, the choke controller adjusts backpressure based on measured wellhead pressure, which itself is influenced by the influx and by the controller's actions. If the estimator ignores this feedback, it may misattribute controller-induced oscillations to influx signatures. The framework addresses this by including a simplified representation of the controller in the input model when its logic is known, or by treating choke position as the effective input and modeling the resulting pressure response [40]. The disturbance-aware likelihood further reduces sensitivity to controller idiosyncrasies by modeling residuals correlated with control increments.

A practical decision output is a recommended diagnostic action rather than a definitive conclusion. For example, when the posterior indicates moderate risk but high uncertainty due to confounding drift, the framework can suggest a small probing choke perturbation that is predicted to yield high information gain. While the present paper does not optimize such perturbations explicitly, the machinery required to evaluate their expected effect is present: one can simulate the predictive distribution under candidate perturbations and compute a scalar measure of posterior contraction. This reframes diagnostic practice as an experiment design problem embedded in routine control, providing a pathway to reduce ambiguous states without resorting immediately to full shut-in procedures.

Another decision output is a consistency score between inferred multiphase state and auxiliary measurements when available [41]. If flow-out and pit volume are measured, the estimator can compute the posterior predictive distribution of these channels and compare with observations. Persistent inconsistency can indicate sensor faults or unmodeled events such as lost circulation. The same Bayesian machinery can be extended to include hypotheses for losses, allowing the system to discriminate between influx and losses when both can produce similar pit-volume trends. The present paper focuses on influx, but the architecture is extensible because it is built around a generative model and explicit priors.

7. Implementation Pathway and Validation Protocol

Deploying the proposed framework requires choices about data handling, parameter initialization, computational platform, and validation methodology. This section outlines a concrete implementation pathway emphasizing technical details that affect estimator stability and interpretability [42].

The data pipeline begins with synchronization and filtering. Surface sensors often have different sampling rates and time stamps. Wellhead pressure may be sampled at high rate, while pit volume may be lower rate and lagged. The estimator operates on a unified time grid at rate Δt. Measurements are aligned to this grid with interpolation where appropriate, but care is taken not to introduce artificial smoothness that would distort innovations. A modest anti-alias filter may be applied to pressure to remove high-frequency noise that is clearly outside the model's bandwidth, but aggressive filtering is avoided because early kick signatures can appear as small slope changes that would be attenuated.

Initialization uses a nominal steady-state profile computed from known mud weight, geometry, and pump rate [43]. However, in practice, nominal profiles are uncertain. The estimator therefore initializes with a prior distribution over reduced coordinates centered on the nominal, with covariance reflecting uncertainty in friction and density. Nuisance parameters such as choke coefficient and pressure sensor bias are initialized with broad priors. The estimator then undergoes a calibration phase during normal

drilling in which it adapts nuisance parameters while keeping influx coefficients near zero due to sparsity priors. This calibration phase is not a separate mode; it is simply the natural behavior of the posterior when the likelihood is consistent with no influx.

Computationally, the reduced-order model and particle filter are well-suited to edge deployment on a rig-site server [44]. The dominant operations are evaluating nonlinear functions and projecting constraints. Efficient linear algebra is used for basis transformations, and the constraint projections are implemented with closed-form clipping and mass-conserving adjustments rather than generic optimization. The particle filter is parallelizable across particles, enabling acceleration on multi-core CPUs. Memory usage is modest because the reduced state dimension is small and only a fixed lag of particles is stored for smoothing.

A key implementation detail is numerical integration. Because the reduced model arises from a projection of conservation laws, it can inherit stiffness from compressibility and relaxation dynamics [45]. The framework uses an implicit-explicit integration strategy: stiff relaxation toward equilibrium for dissolved gas is handled implicitly or with an exact exponential update, while advection-like dynamics are handled explicitly with small Δt. Because the estimator must run at fixed measurement cadence, the integration is tuned so that each time step completes within the available cycle. If the measurement rate is high, sub-stepping may be used within Δt for stability while still providing outputs at each measurement time.

Validation requires both synthetic and experimental components. Synthetic validation uses a higher-fidelity simulator, potentially with finer spatial resolution and more detailed closures, as a ground truth generator. The estimator is then run using only surface pressure and inputs, and its posterior is compared to the known ground truth influx parameters [46]. This assesses identifiability and bias under controlled mismatch. Because the goal is robust inference under mismatch, validation includes stress tests in which the estimator's model intentionally differs from the generator in friction factors, slip behavior, and choke maps. Performance is measured not only by point estimation error but also by calibration of uncertainty: credible intervals should contain true values with appropriate frequency, and alarm probabilities should align with actual events.

Experimental validation can leverage laboratory flow loops and tower facilities that reproduce annular multiphase behavior. A robust protocol includes a range of influx sizes, durations, and fluid types, with both static and circulating conditions, and with controlled choke actuation sequences. The estimator is run on the same surface measurement streams that would be available in field deployment [47]. Performance is evaluated in terms of detection delay, false-positive rate under no-kick scenarios, and the fidelity of inferred migration speed and free gas volume. Importantly, validation should include scenarios with deliberate control perturbations to test whether the disturbance-aware likelihood can separate actuation-induced effects from influx signatures.

Field validation must account for the fact that true influx parameters are rarely known precisely. Therefore, validation focuses on consistency with operational markers, such as known well-control events, and on the estimator's ability to provide early warning relative to existing practices. The estimator's probabilistic outputs facilitate such evaluation: rather than claiming a definitive detection time, one can examine when R_k crosses successive thresholds and how those crossings align with operator decisions and subsequent evidence. When downhole sensors are available, even intermittently, they can be used as partial ground truth to assess inferred pressure profiles and gas-fraction evolution [48]. Another valuable field validation approach is retrospective analysis of historical data with known incidents, using the estimator to reconstruct posterior trajectories and compare against post-event interpretations.

A practical consideration is operator trust. Even a technically correct estimator may be ignored if its outputs are not interpretable. The framework's outputs are designed to be interpretable: posterior distributions over influx mass and time-to-surface can be summarized as intervals, and the estimator can provide diagnostic explanations in terms of which residual components cannot be explained by control-correlated effects. This does not require revealing internal model details, but it does require careful visualization and logging. The estimator can record the innovation decomposition, showing when and

how the residual shifts from control-correlated to unexplained, which can aid human decision-making [49].

8. Discussion

The proposed framework sits between two extremes: high-fidelity simulation and purely data-driven classification. Its purpose is not to replace either, but to provide a real-time inference layer that is fast enough for online use and structured enough to generalize beyond narrow training regimes. Several implications follow.

First, the role of reduced-order modeling is not simply to accelerate simulation. In inversion, reduction also acts as a regularizer: by restricting the state to an observable manifold, it limits the estimator's ability to overfit noise. This is particularly important when only surface pressure is available [50]. However, reduction can also remove critical degrees of freedom if the basis is poorly chosen. Therefore, basis construction should emphasize modes that influence wellhead pressure and that are likely to be excited by influx events. Offline simulation across a broad operating envelope is one way to build such bases, but online adaptation may be needed when well geometry or mud properties differ substantially from prior scenarios.

Second, thermodynamic consistency matters in a diagnostic sense, not merely in prediction accuracy. When dissolution is significant, early influx mass may not appear as free gas, and pressure responses can be altered. Enforcing an equilibrium-constrained dissolved component prevents the estimator from using unphysical free-gas states to match pressure data, which would lead to biased influx size estimates and potentially incorrect control recommendations [51]. At the same time, strict equilibrium may be violated in reality due to kinetics and mixing limitations. The relaxation-based enforcement used here provides a compromise: it encodes equilibrium as an attracting manifold while permitting deviations that can be inferred as part of nuisance behavior.

Third, disturbance-aware likelihood design is essential for operation under managed-pressure control. In practice, many false alarms arise not because influx signatures are weak, but because control actions produce residual patterns that naive models misinterpret. By explicitly modeling residual components correlated with control increments, the estimator becomes more conservative during routine actuation while remaining sensitive to anomalies that do not align with those patterns. This is conceptually similar to including a learned disturbance model, but embedded within a Bayesian estimator so that uncertainty and drift are treated consistently [52].

Fourth, the probabilistic framing changes how alarms can be used. A binary alarm is inherently brittle. A probability enables graded responses, including increased monitoring, diagnostic probing, and escalation to shut-in. It also supports post hoc analysis of system performance through calibration curves: if events with $R_k = 0.7$ occur about 70% of the time, the system is calibrated; if not, priors and noise models can be adjusted. This calibration perspective is valuable because it ties estimator behavior to measurable operational outcomes.

Finally, the framework highlights a path toward active diagnosis [53]. Managed-pressure drilling systems already have actuation authority; using small, safe perturbations to improve identifiability is a natural extension. Because the estimator provides a predictive distribution under candidate actions, one can select actions that are expected to reduce posterior uncertainty about influx parameters. This shifts the paradigm from passive detection to experiment-driven inference, potentially reducing detection delay without increasing false alarms. Implementing active diagnosis requires integration with control safety constraints and human oversight, but the underlying inference machinery is a prerequisite.

Limitations remain. If surface pressure is the only reliable measurement and operations are highly variable, there will be regimes where the inversion remains ambiguous [54]. For example, rapid changes in cuttings loading or unknown losses can produce pressure behaviors that confound influx inference. The Bayesian architecture can incorporate additional hypotheses, but model misspecification remains a risk. Moreover, reduced-order models depend on the representativeness of their bases; if the well configuration changes abruptly, basis mismatch can degrade performance. These limitations motivate

hybrid deployment: the estimator can run alongside existing rule-based systems, providing probabilistic context rather than a single point of failure, and it can be updated as new data becomes available.

9. Conclusion

This paper presented a technical framework for real-time kick diagnosis that treats surface pressure as a control-contaminated observation of a constrained latent multiphase state. The core contribution is a thermodynamically consistent reduced-order model coupled to a sequential Bayesian inference engine with a disturbance-aware likelihood [55]. The reduced-order model preserves the dominant pressure and gas-fraction modes that are observable at the wellhead while enforcing admissibility constraints and an equilibrium-attracting representation of dissolved gas behavior. The inference engine jointly estimates latent state, influx parameters, and slowly drifting nuisance parameters representing friction, choke-map uncertainty, and sensor bias, yielding probabilistic alarms and uncertainty-bounded influx descriptors. A likelihood decomposition that models control-correlated residual components was introduced to reduce false alarms during routine managed-pressure drilling actuation, enabling more robust anomaly attribution to influx dynamics.

The framework outputs not only detection probabilities but also posterior distributions over derived quantities such as influx mass and time-to-surface, supporting risk-ranked operational decisions and control-compatible forecasting. An implementation pathway was outlined, including numerical integration choices, constraint projection for estimator stability, and validation protocols spanning synthetic stress tests, laboratory experiments, and field retrospective analysis. The approach is designed to be extensible: additional hypotheses such as losses or sensor faults can be incorporated within the same generative modeling and Bayesian inference structure. Overall, the results support the thesis that uncertainty-aware inversion built on physics-guided reduced-order modeling can provide actionable, interpretable, and computationally feasible kick diagnosis using surface pressure data in managed-pressure drilling environments [56].

References

[1] K. M. Suyan, D. Dasgupta, D. Sanyal, V. Sharma, and V. K. Jain, "Managing total circulation losses with crossflow while drilling: Case history of practical solutions," in *SPE Annual Technical Conference and Exhibition*, SPE, 11 2007.

[2] 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.

[3] M. Alsheikh, C. Gooneratne, A. Magana-Mora, M. Ibrahim, M. Affleck, W. Contreras, G. D. Zhan, M. A. Jamea, I. A. Umairin, A. Zaghary, M. ilies Ayachi, A. G. Abdel-Kader, S. Ahmed, G. Makowski, and H. Kapoor, "Internet of things iot edge computer vision systems on drilling rigs," in *SPE Middle East Oil & Gas Show and Conference*, SPE, 12 2021.

[4] G. J. Moridis, T. S. Collett, M. Pooladi-Darvish, S. H. Hancock, C. Santamarina, R. Boswell, T. J. Kneafsey, J. Rutqvist, M. B. Kowalsky, M. T. Reagan, E. D. Sloan, A. K. Sum, and C. A. Koh, "Challenges, uncertainties, and issues facing gas production from gas-hydrate deposits," *SPE Reservoir Evaluation & Engineering*, vol. 14, pp. 76–112, 2 2011.

[5] D. A. Elliott, "Evaluation of results from one company's global use of underbalanced drilling for tight gas exploration and appraisal reservoir characterization," in *SPE/IADC Managed Pressure Drilling and Underbalanced Operations Conference and Exhibition*, SPE, 3 2012.

[6] T. R. Salakhov, V. Dubinsky, and V. U. Yamaliev, "A field-proven methodology for real-time drill bit condition assessment and drilling performance optimization," in *SPE Russian Oil and Gas Technical Conference and Exhibition*, SPE, 10 2008.

[7] M. Okot, M. Campos, G. Muñoz, A. Alalsayednassir, M. Weber, and Z. Muneer, "Utilization of an innovative tool to improve hole cleaning efficiency in extended reach wells in saudi arabia," in *SPE Kuwait Oil and Gas Show and Conference*, SPE, 10 2015.

[8] I. Kocis, T. Kristofic, M. Gajdos, G. Horvath, and S. Jankovic, "Utilization of electrical plasma for hard rock drilling and casing milling," *SPE/IADC Drilling Conference and Exhibition*, 3 2015.

[9] 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.

[10] Y. Lu, Z. Zhou, Z. Ge, X. Zhang, and Q. Li, "Research on and design of a self-propelled nozzle for the tree-type drilling technique in underground coal mines," *Energies*, vol. 8, pp. 14260–14271, 12 2015.

[11] R. L. Spross, "Halliburton sperry-sun doe high temperature lwd project," 3 2005.

[12] M. P. Starrett, A. D. Hill, and K. Sepehrnoori, "A shallow-gas-kick simulator including diverter performance," *SPE Drilling Engineering*, vol. 5, pp. 79–85, 3 1990.

[13] 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.

[14] "Basic data report for drillhole erda 6 (waste isolation pilot plant - wipp)," 1 1983.

[15] J. T. Finger and B. J. Livesay, "Alternative wellbore lining methods: Problems and possibilities," 8 2002.

[16] F. Berge, N. Hepburn, B. J. S, and T. Gronas, "A new generation multilateral system for the troll olje field - development and case history," *SPE Offshore Europe Oil and Gas Exhibition and Conference*, 9 2001.

[17] Y. Zhang, W. Wang, P. Zhang, G. Li, S. Tian, J. Lu, and B. Zhang, "A solution to sand production from natural gas hydrate deposits with radial wells: Combined gravel packing and sand screen," *Journal of Marine Science and Engineering*, vol. 10, pp. 71–71, 1 2022.

[18] C. John, G. Maciasz, and B. Harder, "Gulf coast geopressured-geothermal program summary report compilation. volume 2-a: Resource description, program history, wells tested, university and company based research, site restoration," 6 1998.

[19] J. Labaj and D. Barta, "Unsteady flow simulation and combustion of ethanol in diesel engines," *Communications - Scientific letters of the University of Zilina*, vol. 8, pp. 27–37, 6 2006.

[20] A. R. Brandt, "Oil depletion and the energy efficiency of oil production: The case of california," *Sustainability*, vol. 3, pp. 1833–1854, 10 2011.

[21] C. Wei, Y. Chen, O. Santos, M. Kunju, S. Mahmud, M. Almeida, and P. Sonnemann, "An evaluation of pressure control methods during riser gas handling with mpd equipment based on transient multiphase flow modeling and distributed fiber optic sensing," in *SPE Annual Technical Conference and Exhibition*, SPE, 9 2022.

[22] S. M. M. Sarshar, "The applications of a novel compact separation system in ubd and mpd operations," in *IADC/SPE Managed Pressure Drilling and Underbalanced Operations Conference and Exhibition*, SPE, 4 2013.

[23] G. Mouillerat and F. Silva, "Clov project: Overview," in *Offshore Technology Conference*, OTC, 5 2015.

[24] X. J. Zhang, S. Taoutaou, Y. Guo, Y. L. An, S. M. Zhong, and Y. Wang, "Engineering cementing solution for hutubi underground gas storage project," in *SPE Unconventional Gas Conference and Exhibition*, SPE, 1 2013.

[25] R. C. Thompson, B. Murphy, K. Williams, H. S. Giberson, J. Garaghty, and G. Pace, "Energy storage for natural gas fueled electric drilling rigs," *SPE/IADC Drilling Conference*, 3 2013.

[26] T. C. Chidsey, "Production analysis: Cherokee and bug fields, san juan county, utah," 12 2003.

[27] J. H. Hallman, M. A. Muqeem, I. Cook, H. A. Shammari, and C. M. Jarrett, "Fluid customization and equipment optimization enable safe and successful underbalanced drilling of high-h2s horizontal wells in saudi arabia," in *IADC/SPE Managed Pressure Drilling & Underbalanced Operations*, SPE, 3 2007.

[28] 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.

[29] F. Loux, M. J. Bunyak, N. Kongpat, K. Pattanapong, P. Eagark, I. Rubianto, F. Gallo, J. Intravichit, K. Charnvit, A. E. Prasetia, and N. H. Houng, "Managed pressure drilling significantly reduces drilling costs on extremely challenging ultraht wells in the gulf of thailand," in *SPE/IATMI Asia Pacific Oil & Gas Conference and Exhibition*, SPE, 10 2015.

[30] "Two-phase reaction turbine. technical progress report for the period july-december 1999," 12 1999.

[31] N. Gruber and K. L. Adair, "New laboratory procedures for evaluation of drilling induced formation damage and horizontal well performance," *Journal of Canadian Petroleum Technology*, vol. 34, 5 1995.

[32] J. Fiore, "Site-specific development plan: Carlin, nevada," 1 1980.

[33] E. D. Gomez, M. Chavarria, J. C. Beltran, C. P. M. Lupo, H. Duno, and J. Casaos, "Successful use of managed pressure drilling (mpd) in low pressure, high pressure, and deeper reservoirs in mexico south," in *SPE/IADC Managed Pressure Drilling and Underbalanced Operations Conference and Exhibition*, SPE, 3 2012.

[34] "Western gas sands porject. status report, 1 november 1979-30 november 1979," 1 1979.

[35] E. Leusheva, N. Alikhanov, and N. Brovkina, "Study on the rheological properties of barite-free drilling mud with high density," *Journal of Mining Institute*, vol. 258, pp. 976–985, 12 2022.

[36] Y. Lu, S. Xiao, Z. Ge, Z. Zhou, and K. Deng, "Rock-breaking properties of multi-nozzle bits for tree-type drilling in underground coal mines," *Energies*, vol. 9, pp. 249–, 3 2016.

[37] H. N. Patel, J. Bruton, E. Dietrich, and R. Prince-Wright, "Lessons learned and safety considerations for installation and operation of a managed pressure drilling system on classed floating drilling rigs," in *Offshore Technology Conference*, OTC, 5 2015.

[38] "The objectives for deep scientific drilling in yellowstone national park," 1 1987.

[39] A. Meziou, M. Chaari, M. A. Franchek, K. M. Grigoriadis, R. Tafreshi, and B. Ebrahimi, "Subsea production two-phase flow modeling and control of pipeline and manifold assemblies," in *Volume 1: Active Control of Aerospace Structure; Motion Control; Aerospace Control; Assistive Robotic Systems; Bio-Inspired Systems; Biomedical/Bioengineering Applications; Building Energy Systems; Condition Based Monitoring; Control Design for Drilling Automation; Control of Ground Vehicles, Manipulators, Mechatronic Systems; Controls for Manufacturing; Distributed Control; Dynamic Modeling for Vehicle Systems; Dynamics and Control of Mobile and Locomotion Robots; Electrochemical Energy Systems*, American Society of Mechanical Engineers, 10 2014.

[40] J.-M. Sur, H.-N. Im, I.-S. Lee, H. Lee, and J.-W. Choi, "Study on design optimization of degasser baffles using cfd," *Journal of Ocean Engineering and Technology*, vol. 29, pp. 331–341, 10 2015.

[41] F. Deng, H. Liu, S. Chen, G. Chen, M. Wang, C. Liu, J. Shi, and W. Jiang, "Low-cost multi-phase flow metering and assays technology using online magnetic resonance," in *ADIPEC*, SPE, 10 2022.

[42] C. Stone, "Geothermal energy abstract sets. special report no. 14," 1 1985.

[43] E. G. Leonov and V. I. Isaev, *Applied Hydro-Aeromechanics in Oil and Gas Drilling - Stationary Flow of Gas and Gas-Cutting Mixture in Elements of Well Circulation System*. Wiley, 10 2009.

[44] J. H. Cohen, C. E. Leitko, and R. C. Long, "High-power slim-hole drilling system," in *IADC/SPE Drilling Conference*, SPE, 2 2000.

[45] J. Zhang and D. Li, "Research on optimization of large flow hydraulic punching process in zhaozhuang mine," *Geofluids*, vol. 2022, pp. 1–14, 12 2022.

[46] Q. Wang, Z. Wang, J. Yue, F. An, J. Dong, and W. Ke, "Temperature of the core tube wall during coring in coal seam: Experiment and modeling.," *ACS omega*, vol. 7, pp. 7901–7911, 2 2022.

[47] "The yucca mountain project prototype testing program; 1989 status report," 10 1989.

[48] L. Kozlov, Y. Buriennikov, O. Rusu, V. Pyliavets, V. Kovalchuk, O. Petrov, and I. Rusu, "Algorithm of controlling an adaptive hydraulic circuit for mobile machines," *International Journal of Modern Manufacturing Technologies*, vol. 13, pp. 79–86, 12 2021.

[49] D. N. Rao, "Development and optimization of gas-assisted gravity drainage (gagd) process for improved light oil recovery," 10 2003.

[50] Y. Tang, P. Zhao, X. Li, X. Fang, and P. Yang, "Numerical simulation and experimental test of the sliding core dynamics of a pressure controlled jet crushing tool for natural gas hydrate exploitation," *Processes*, vol. 10, pp. 1033–1033, 5 2022.

[51] J. J. Lewis, "New uses of existing technology for controlling blowouts: Chronology of a blowout offshore louisiana," *Journal of Petroleum Technology*, vol. 30, pp. 1473–1480, 10 1978.

[52] H. Nickens, "A dynamic computer model of a kicking well," *SPE Drilling Engineering*, vol. 2, pp. 159–173, 6 1987.

[53] "Assessment of costs and benefits of flexible and alternative fuel use in the us transportation sector," 7 1991.

[54] L. Umar, I. A. B. A. Aziz, N. A. Nordin, R. M. Ali, A. Waguih, F. Rojas, S. Fey, B. Subroto, B. Dow, and G. I. Garcia, "Demonstrating the value of integrating fpwd measurements with managed pressure drilling to safely drill narrow mud weight windows in hp/ht environment," in *SPE/IADC Managed Pressure Drilling and Underbalanced Operations Conference and Exhibition*, SPE, 3 2012.

[55] A. Willersrud, L. Imsland, M. Blanke, and A. Pavlov, "Early detection and localization of downhole incidents in managed pressure drilling," in *SPE/IADC Managed Pressure Drilling and Underbalanced Operations Conference & Exhibition*, SPE, 4 2015.

[56] Y. Zhang, F. Ji, and Q. Zou, "Coordinated slag disposal from horizontal boreholes during hydraulic cutting based on two-phase flow theory," *Frontiers in Earth Science*, vol. 10, 4 2022.