AUTONOMOUS LIQUID UNLOADING OPTIMIZATION FOR INTERMITTENT GAS WELLS USING AN EDGE-DEPLOYED IIOT CONTROL SYSTEM

R. Marin Nunez, A. Gambaretto, C. Kemp, M. Perezhogina, V. Er, D. Davalos, G. Martinez Loya, M. Nethi, V. Salvi, E. Gies

SLB

ABSTRACT

Intermittent gas wells in high-pressure basins face persistent production losses from liquid loading and the limitations of manual or calendar-driven control schemes. This paper describes an autonomous, edge-deployed control application developed to address these challenges in the Haynesville Basin.

The system runs on a local IIoT gateway installed at the wellsite. It ingests real-time casing and tubing pressures, flow rates, and temperature to evaluate well conditions continuously. Using embedded flow physics and a lightweight machine learning model, the system decides when to shut in the well and how long to keep it closed, then reopens the choke at the calculated optimal moment — all without operator input or cloud dependency for real-time decisions.

A nine-well pilot across eight pads demonstrated consistent production improvement. Over periods of 43 to 83 days, cumulative gas output increased between 17% and 139% relative to pre-deployment baselines. Daily rate gains reached up to 350 MCFD per well. Projected annual benefit per well exceeds 80 MMCF of additional gas, valued at over $135,000, against a yearly operating cost of roughly $3,600 per well.

The results confirm that edge-based autonomous control can modernize intermittent well management at low capital cost, using infrastructure already present at most wellsites.

INTRODUCTION

Intermittent production is one of the most common — and least optimized — operating modes in mature gas basins. In the Haynesville and similar high-pressure formations, wells that once flowed continuously begin to cycle as reservoir pressure declines. Bottomhole pressure eventually falls below the threshold needed to sustain continuous liquid transport, and the well must be periodically shut in to allow pressure to rebuild before the next unloading cycle. This behavior places these assets squarely within the domain of artificial lift management, even when no downhole equipment is involved: the cycle timing, choke control, and pressure recovery strategy all directly determine production efficiency.

Conventional management of intermittent wells relies on fixed schedules or simple pressure ratio triggers programmed into SCADA systems. Both approaches are static — they cannot respond to the natural variability in reservoir behavior, wellbore geometry, or liquid loading severity from one cycle to the next. The result is a combination of premature shut-ins that waste production time, late shut-ins that allow deep liquid loading and extend recovery time, and frequent field visits to manually adjust choke settings.

This paper presents a field-deployed solution that replaces static control with a continuous, condition-aware system. Running on a ruggedized gateway at the wellsite, the application evaluates real-time sensor data against physics-based flow models to detect the onset of loading and initiate shut-in at the right moment. Reopening timing is guided by a machine learning model trained on that well's own pressure buildup and flowback history. The combination eliminates the need for operator scheduling while adapting to each well's unique behavior.

From a commercial standpoint, this approach targets the economics of late-life assets specifically. Intermittent wells typically operate under constrained budgets, which rules out expensive sensor upgrades or custom automation builds. By leveraging the instrumentation already present at most Haynesville pads — casing and tubing pressure gauges, a flowmeter, and a remote-actuated choke — the system delivers optimization with minimal incremental capital.

METHODOLOGY

System Overview

The control system is a containerized software application deployed on an IIoT field gateway. It connects to existing wellhead sensors and the choke actuator through the site's RTU, requiring no changes to SCADA configuration or communications infrastructure. All real-time decisions are made locally on the gateway device.

The application operates through three functional stages that run continuously in a closed loop:

Flow Condition Assessment: Sensor telemetry — casing pressure, tubing pressure, temperature, and surface flowrate — is ingested and validated at the edge. From this data, the system computes two key indicators in real time. Gas velocity is derived from surface flowrate and tubing geometry. Critical velocity is calculated using Turner's mechanistic model for multiphase flow, defining the minimum gas rate needed to carry liquids to surface. Liquid column height is inferred from the differential between casing and tubing pressures, corrected for gas compressibility and temperature. The Coleman, Nosseir, and Barnea models are also supported for operators who prefer alternative flow regime assumptions.

Shut-In Decision: The control algorithm monitors gas velocity relative to the critical velocity threshold and tracks the rate of liquid accumulation inferred from pressure differentials. When gas velocity falls below the critical threshold and the estimated liquid column approaches a calculated tipping point, the system closes the choke. Critically, the algorithm distinguishes short-lived flow fluctuations from sustained degradation by requiring consistency across multiple time windows before acting — avoiding unnecessary shut-ins while still catching genuine loading events before they become severe.

Reopening Optimization: Rather than applying a fixed shut-in duration, the system uses a machine learning model to predict when pressure has recovered sufficiently to support effective unloading. The model outputs a recommended wait time based on that well's current pressure buildup trajectory and prior cycle performance. The choke reopens automatically at the predicted optimum.

Figure 1 - Shut-in trigger logic: real gas rate versus critical gas rate with liquid loading onset

Machine Learning Model for Shut-In Duration

The ML component addresses a fundamental limitation of fixed-timer approaches: pressure recovery speed varies well by well and cycle by cycle, depending on reservoir behavior, wellbore temperature, and the degree of loading at shut-in. A single recovery time cannot be optimal across all conditions.

The workflow uses cycle-segmented data — each cycle defined as the shut-in phase plus the subsequent flow period. Segmenting this way isolates pressure buildup behavior from the noise of active flow and improves model stability. Only cycles with confirmed successful unloading are included in training, preventing the model from learning from failed or anomalous events.

Two models work in sequence. The first is trained on pressure measurements recorded immediately after shut-in and learns to extrapolate the full pressure buildup curve from the early trend. The second model takes the projected buildup curve and prior cycle statistics as inputs and predicts average production rate as a function of proposed shut-in duration. Plotting that relationship reveals a clear optimum — the duration beyond which additional wait time yields no further production benefit. The system targets that point for each cycle.

Both models are maintained independently for each well. They are retrained automatically as new cycles accumulate, allowing the system to adapt to gradual changes in reservoir pressure, liquid loading patterns, and equipment behavior without manual intervention.

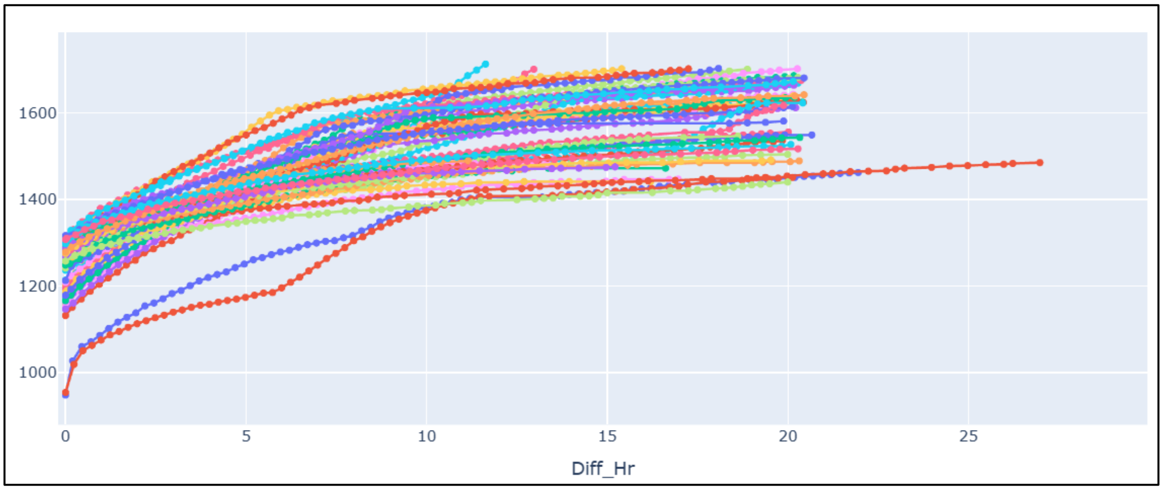

Figure 2 - Tubing head pressure buildup with noise removed

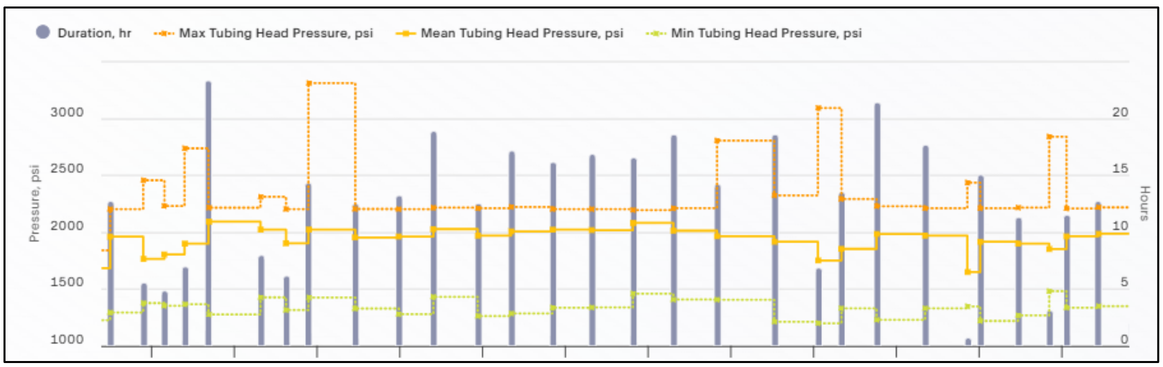

Figure 3 - Shut-in phase cycle statistics over time

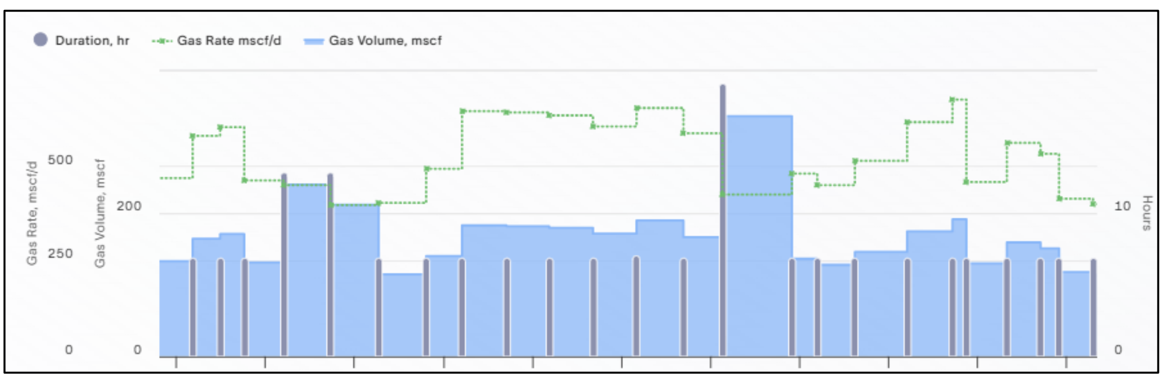

Figure 4 - Flow phase cycle statistics over time

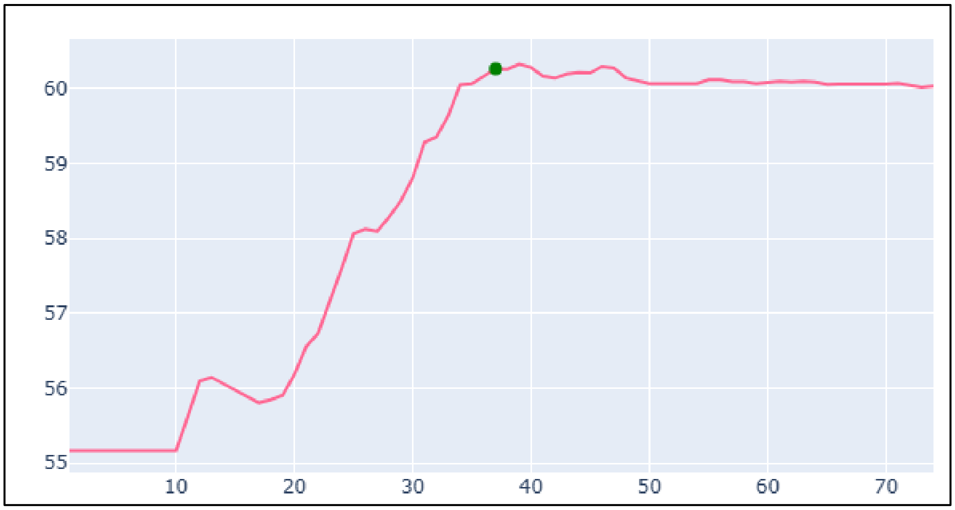

Figure 5 - Predicted production rate versus shut-in duration with optimal point identified

Cloud Layer and ML Operations

Real-time well control operates entirely at the edge and does not require cloud connectivity. However, the ML models are hosted and retrained in the cloud to simplify lifecycle management across a distributed fleet of devices.

When a shut-in begins, the edge device captures the initial pressure buildup data and transmits a compact cycle summary to the cloud. The cloud-hosted model processes this summary and returns the recommended shut-in duration within seconds. The edge device uses this value to schedule the reopen command. Full-resolution time-series data is retained locally; only summarized cycle metadata is transmitted, keeping bandwidth consumption low — an important consideration for sites on cellular or satellite connections.

Retraining is triggered automatically as new cycle data accumulates. Model updates are deployed back to the inference service without requiring firmware changes or physical access to field devices. This centralized approach makes it practical to maintain well-specific models across large pad counts with minimal engineering overhead.

Operational Safeguards

Fully autonomous operation requires reliable mechanisms to detect and escalate equipment or flow anomalies that the control algorithm cannot resolve on its own. Three safeguard conditions are actively monitored. A non-response condition is flagged when the choke fails to confirm a position change within the expected window, which may indicate a hardware fault or communication failure. A leakage condition is flagged when flow is measured while the choke is commanded closed. A no-flow condition is flagged when the choke is open but the well is not producing, which may indicate severe backpressure, a failed choke, or a downhole obstruction.

When any of these conditions is detected, an alert is forwarded to the cloud platform where the responsible engineer receives a notification. This management-by-exception approach means that engineering attention is directed only to wells that actually need it, rather than spread thin across routine monitoring of every asset in the fleet.

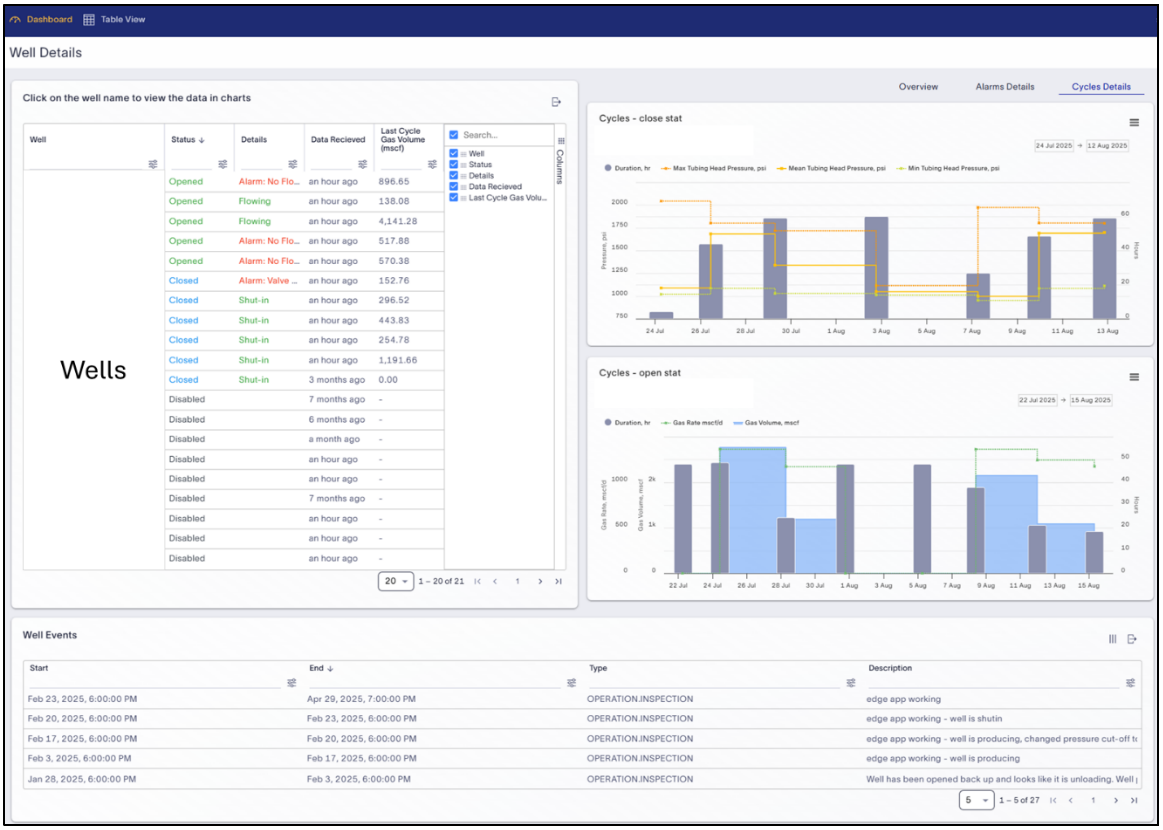

Figure 6 - Fleet monitoring dashboard

Figure 7 - Exception alert queue

Per-Well Configuration

The system supports per-well operating constraints configured through the cloud interface. These include minimum and maximum allowable shut-in durations, minimum and maximum open durations, production rate thresholds for triggering closure, and fallback modes that revert to pressure-ratio logic or fixed cycling if sensor data quality degrades. These boundaries prevent the control algorithm from acting on outlier model outputs or anomalous sensor readings while preserving the core autonomous behavior under normal conditions.

FIELD IMPLEMENTATION AND RESULTS

Pilot Scope and Well Selection

The pilot covered nine wells across eight pads in the Haynesville Basin. The well population was intentionally heterogeneous: some wells cycled slowly with long shut-in periods and rapid flowback, while others cycled frequently with shorter recovery windows and moderate loading. This range was chosen to test whether a single control architecture could perform across meaningfully different operating conditions.

Well selection was based entirely on existing infrastructure. Each candidate well required an RTU with communication capability, a remote-actuated choke, reliable cellular coverage, and available solar power. No new sensors or mechanical modifications were installed. The application was deployed as a containerized package on IIoT gateways at each pad, with multiple wells grouped under a single device where pad layout allowed. Installation required only software configuration and sensor mapping through the existing RTU.

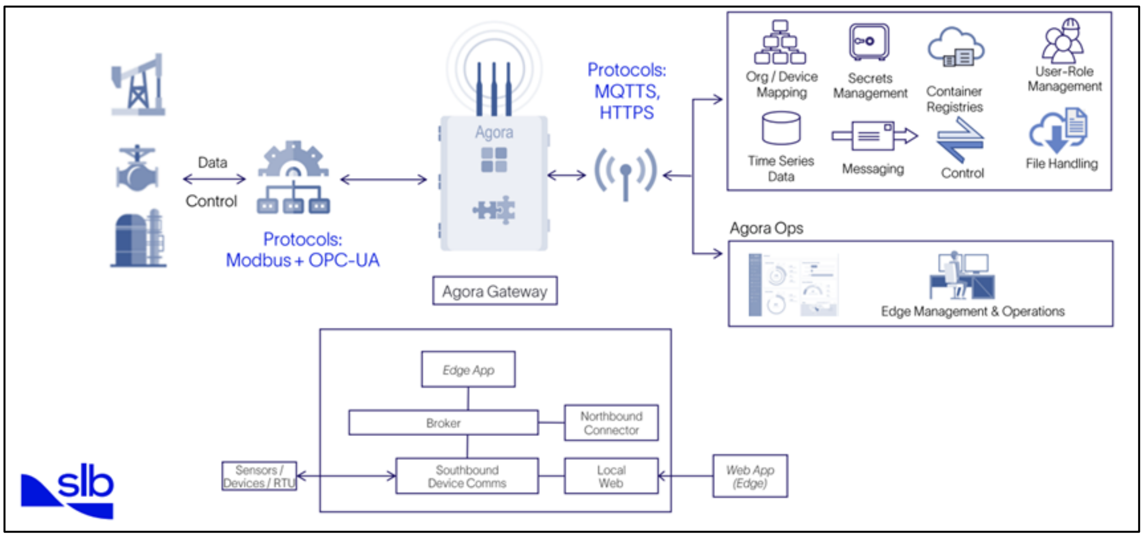

Figure 8 - Pad-level edge deployment architecture

Production Results

The system operated autonomously across all nine wells for periods ranging from 43 to 83 days. Production was compared against a capacity baseline derived from each well's pre-deployment decline trend. Results are summarized in Table 1 below.

Table 1 - Cumulative Production Results by Well

Well	Days	Baseline Cum. (MCF)	Optimized Cum. (MCF)	Change
Well 1	63	5,864	14,013	+139.0%

Well 2	83	10,402	22,338	+114.8%
Well 3	83	14,534	24,808	+70.7%
Well 4	43	4,871	7,999	+64.2%
Well 5	43	7,412	10,885	+46.8%
Well 6	48	3,939	5,698	+44.7%
Well 7	43	8,914	12,214	+37.0%
Well 8	83	22,026	27,341	+24.1%
Well 9	63	9,785	11,444	+16.9%

Aggregate outcomes across the nine-well pilot include cumulative production gains ranging from 17% to 139%, daily rate improvements of up to 350 MCFD, and projected annual incremental production exceeding 80 MMCF per well. The economic case is straightforward: at prevailing gas prices, that volume represents over $135,000 in annual value per well against a system operating cost of approximately $3,600 per year.

Notably, field personnel were not yet fully familiar with autonomous operations during this initial pilot, and some manual well checks were still performed. Even under these conditions, the system met its production objectives while eliminating scheduled choke cycling and manual shut-in decisions by operators.

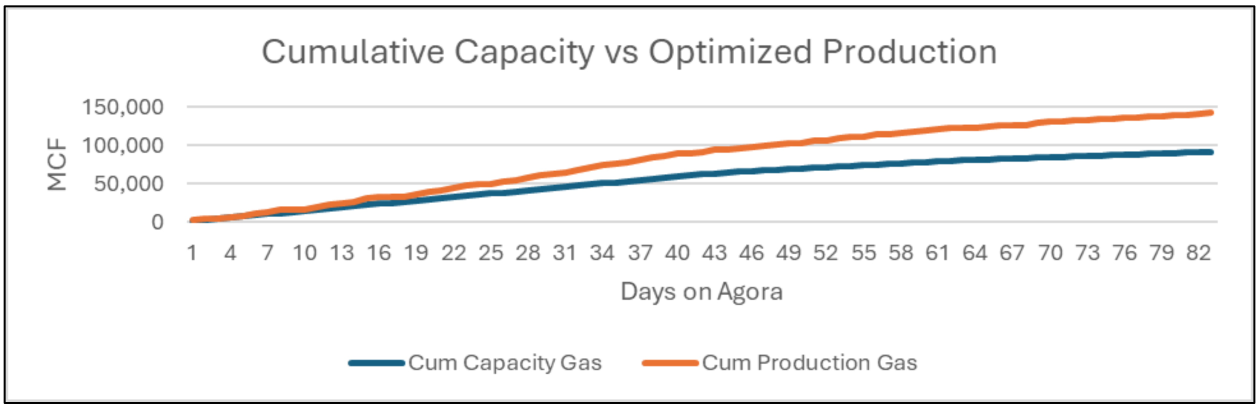

Figure 9 - Baseline versus optimized cumulative production across the well set

Well-Level Analysis

Well 1 produced the largest cumulative gain in the pilot at 139% over 63 days, averaging 225 MCFD above the pre-deployment baseline. The ML model recommended longer shut-in durations than had been used manually, which reduced cycle frequency and mechanical actuation while keeping gas rates above critical velocity for longer stretches of each flow period. The extended flow windows before loading onset were the primary driver of the production gain.

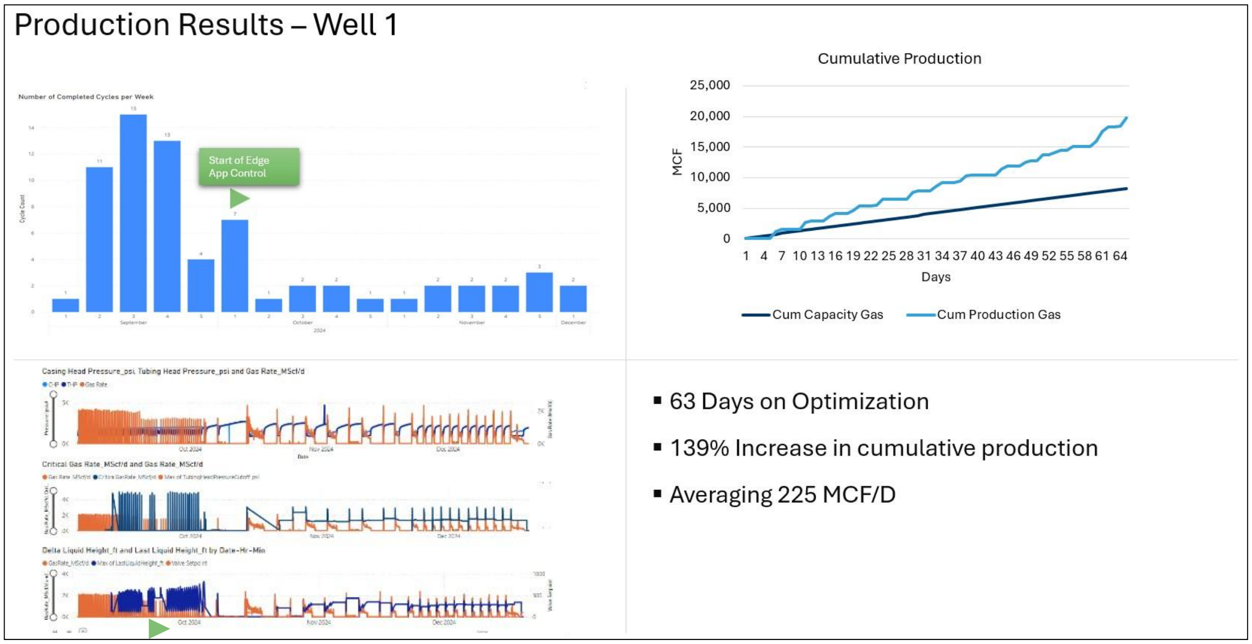

Figure 10 - Well 1 production summary

Well 2 achieved 114% cumulative improvement over 83 days, with an average daily gain of 269 MCFD. A consistent pattern emerged in the cycle-level data: more than half of total gas produced in each cycle occurred within the first third of the open period, with production tailing off sharply thereafter. The algorithm's early shut-in enforcement prevented the well from spending extended time in low-productivity tail flow, which proved to be one of the more impactful efficiency gains across the pilot.

Figure 11 - Well 2 production summary

Well 3 delivered 70% cumulative improvement over 83 days and the highest single-well daily rate uplift in the pilot at 350 MCFD average. The same front-loaded cycle pattern observed in Well 2 was present here, reinforcing the value of timely shut-in across wells with different absolute production rates. ML-driven pressure recharge optimization sustained consistent cycle performance throughout the 83-day window with no manual adjustments.

Figure 12 - Well 3 production summary

Wells 4 through 9 showed improvements ranging from 17% to 65%. While these gains are more modest, they reflect wells with shorter optimization windows (as few as 43 days) or wells that were already cycling near a reasonable baseline. Across all nine wells, no well showed a production decline under autonomous control, confirming that the system's safeguards and fallback logic functioned as intended even on less responsive assets.

CONCLUSION

This paper presents field results from a pilot deployment of an autonomous liquid unloading control system on nine intermittent gas wells in the Haynesville Basin. The system combines real-time flow physics with per-well machine learning to continuously evaluate unloading conditions and manage choke actuation without operator scheduling or cloud dependency for control decisions.

The pilot demonstrated production improvements across all nine wells, with cumulative gains between 17% and 139% and daily rate uplifts reaching 350 MCFD. A consistent finding across the higher-performing wells was the value of early, condition-based shut-in: preventing tail-flow cycles and enabling adequate pressure recharge delivered more production benefit than any single parameter adjustment.

From a deployment standpoint, the solution requires no new sensors or mechanical modifications — only a compatible RTU, a remote-actuated choke, and cellular connectivity. This makes it applicable to a broad population of late-life Haynesville wells and similar assets in other mature basins. The containerized architecture also supports multi-well grouping on a single gateway, keeping per-well hardware costs low.

As well-specific models accumulate more cycle data, prediction accuracy improves and the system adapts to gradual changes in reservoir behavior and equipment condition. This continuous learning capability positions edge-based autonomous control as a practical long-term operating strategy, not just a short-term optimization campaign.

REFERENCES

Chen, Z. et al., 2024. Smart Needle Valve: Intelligent and Cost-Effective Solution to Unlock Gas Production in a Tight Gas Field in China. OTC.

Gambaretto, A. & Kashif, R., 2023. A Robust Method for Data-Driven Gas-Lift Optimization. SPE.

Gambaretto, A., Yermekova, M., Srivastava, S. & Hyder, Z., 2024. Autonomous SRP Optimization Using an Edge-Based Data-Driven Solution. SPE.