Predicting Natural Gas Pipeline Alarms

Colin O. Quinn, Richard J. Povinelli, Jaired R. Collins, and Jake Norment

This paper provides a method to forecast natural gas pipeline pressure alarms. Forecasting natural gas pipeline pressure alarms help control room operators maintain a functioning pipeline and avoid costly down time.

Natural gas production companies use pipelines to transport natural gas from the extraction well to a distribution point. As gas enters the pipeline and travels to the distribution point, it is expected that the gas meets certain specifications set in place by either state law or the customer receiving the gas. If the gas meets these standards and is accepted at the distribution point, the pipeline is referred to as being in a steady-state. If the gas does not meet these standards, the production company runs the risk of being shut-in, or being unable to flow any more gas through the distribution point until the poor-quality gas is removed. Being shut-in is costly and time consuming for the production company. For the pipeline to become functional again, the unexpected gas must either be diffused with gas further down the line or flared from the system entirely.

The internal pressure of the pipeline measured at the distribution point is the variable forecasted to help controllers deliver acceptable gas in a safe and reliable way. Sensors are used to collect real-time pressure information from within the pipe, and alarms are used to alert the control operators when a threshold is exceeded. If operators fail to keep the pipeline's pressure within an acceptable range, the company risks being shut-in or rupturing the pipeline. Predicting pressure alarms enables operators to take appropriate action earlier to avoid being shut-in and is a form of predictive maintenance.

We forecast alarms by using an autoregressive model (AR) in conjunction with alarm thresholds. The alarm thresholds are defined by the production company and are occasionally adjusted to meet current environment conditions. A regression-based approach to predicting alarms is used in favor of a classification-based approach because the alarm thresholds can be changed after the algorithm is deployed. The benefit of a regression-based model is in its output, since it can be used to diagnose the state of the pipeline rather than just classifying an alarm being imminent.

The results of the alarm forecasting method show that we can accurately forecast the pressure time series up to a 30-minute time horizon. This translates into true positive rates that drop of linearly from around 100% at one minute to approximately 65% at a 30-minute forecast horizon. This means that at 30 minutes, we correctly forecast 65% of the alarms, but have no false positives.