Abstract

For decades sucker rod pump artificially lifted wells have used devices called pump off controllers (POC) to match the pumping unit's runtime to the available reservoir production by idling the well for a set time where variable frequencies drives are not available. In doing this the POC allows the well to enter a set period of downtime when the downhole pump fillage is incomplete to avoid premature failures, and then brings the well back online to operate before production is lost. Although this method has been successful for several years, autonomous control algorithms can be utilized to reduce failures or increase production in cases where the downtime is not already optimized. Optimizing the idle time for a sucker rod pump artificially lifted well involves understanding the amount of time required to fill the near wellbore storage area before generating a fluid column above the pump intake that will begin to hinder inflow from the reservoir into the wellbore. By varying the idle time and observing the impact on production and cycles the program hunts for the optimal idle time. By constantly hunting for the optimal idle time the optimization process can adjust the idle time when operating conditions change. This gives the advantage of always meeting the current well bore and reservoir conditions without having to have a user make these changes and determine what the downtime for the well is. Autonomously modulating the idle time for a well, if done properly will either reduces incomplete fillage pump strokes, in cases where the idle time is too short, or will increase the wells production in cases where the idle time is too long. Overall this will result in the optimization of wells by reducing failures and/or increasing production, generating a huge value to the end user by automating the entire process of downtime optimization.

Introduction

Matching a well's production to the artificial lift method's capacity has been a challenge that has needed an innovative solution for many years. One of the most popular forms of artificial lift around the globe is sucker rod pump. One of the primary advatnages of having a sucker rod pump artificial lift system is that the well can achieve maximum drawdawn. This means that the pumping unit can bring the casing fluid level down to the pump intake, which creates the lowest back pressure on the reservoir, allowing the maximum inflow from the reservoir to the well bore. Although this is an advantage of utilizing a sucker rod pump, operating with a fluid level at the pump intake results in incomplete pump fillage often referred to as fluid pound. Fluid pound is well known to be damaging to the downhole pump equipment because it creates sudden and sharp collision between the downhole pump that compresses the rods above the pump as well as disrupts the pump's downhole travel.

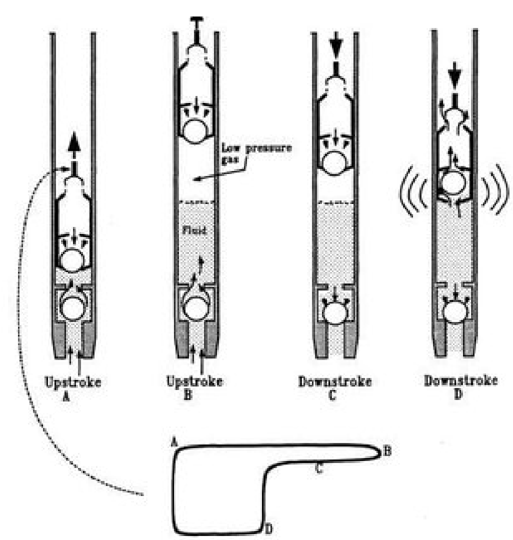

Figure 1: Downhole dynamometer and pump illustration example of fluid pound

The initial solution was to run the well with fluid pound constantly and deal with the consequences of damaging the downhole equipment. After many years of high failure rates, technology was created to improve the runtime of the artificial lift system base on the well's production using a device called a time clock. These were crude devices that allow the wells to turn on and off based on a preconfigured runtime of the well. Although these were improvements in some ways from the basic 24-hour runtime of a well, it was very difficult to optimize these devices as they were controlled by a manual time input and not any real observable data. With the invention of downhole dynamometer cards came modern pump off control systems that offered a way to control the well based on pump fillage. This gave operators a huge advantage because much of the guess work was taken out of operating their wells. They would simply set up their pump off controller to idle the well for a certain amount of time when the controller detects incomplete pump fillage. This allows the pumping unit to completely pump off the fluid level in the casing anulus without running 24 hours a day with incomplete pump fillage. This drastically reduced the fluid pounds strokes in a pumping system without losing production. However, there was one variable left for operators to optimize. The amount of time the well would be "idle" before it came back online and started pumping. If the unit was idle for too long the fluid level in the casing would build up and restrict inflow, harming production. If the unit was idle for too short the well could cycle many times a day generating many fluid pound strokes every time the well pumped off. This period of time the well is idle is often referred to as the well's idle time or downtime.

To optimize the well's downtime, algorithms have been developed to understand the well's operating condition and allows for software programs to hunt for the optimal downtime. By modulating the idle time every 24 hours and observing the effects on production, runtime, and cycles, the software program can determine if the idle time should be increased or decreased. Constantly changing the idle time and observing the impact of these changes allows the software program to accommodate the well's production capacity in real time. This avoids two suboptimal scenarios: scenario 1 where the pumping unit's idle time is too long, and the well is losing production and scenario 2 where the pumping unit's idle time is too short, and the well is experiencing high cycles and unnecessary fluid pound strokes which will lead to premature failures. This solution allows operators to get the most out of their pump off controllers without requiring them to manually optimize their setpoints every day. This ensures that their rod pumping systems are operating according to best practices without requiring the extensive man hours necessary for optimizing these setpoints, allowing them to get the most out of their automation platforms.

Statement of Theory and Definitions

The primary purpose of a rod pump controller is provide a solution that allows a well operating at a fixed speed, meaning a variable frequency drive is not being utilized, to meet the inflow of the reservoir to achieve maximum drawdown for the well without having to run the well in a constant run mode. This is achieved by putting a well in an idle state when the pump fillage begins to decrease. Which means the well will turn off for a prescribed period of time, the downtime, and then return to pumping. Each one of these on-off-on sequences is referred to as a cycle. By idling the well and allowing enough fluid to flow into the volum near the well bore and pump, the rod pump controller allows the well to produce with a full pump and elimate most of the harmful pump strokes that occur when the pump fillage is incomplete. However, if the idle time is too short the well will return to production too soon and quickly pump all of the nearby fluid and quickly return to the incomplete fillage where the well will idle again. Every time the well cycles it requires incomplete fillage to do so, which is difficult on the equipment, so the fewer times the well cycles, the better it will be for the overall health of the rod pumping system. It is also suboptimal to allow the well to idle too long. Although idling the well longer will reduce the cycles and therefor reduce the quantity of incomplete fillage strokes, it also leaves the well velnurable to having restricted inflow due to hydrostatic head from a fluid column building up above the pump. The fluid level above pump causes a pressure against the reservoir, which will eventually reach equilibrium once the pressure from the column of fluid in the casing anulus equals the pressure caused from the reservoir.

Optimally the well's idle time should be short enough where the well is not allowing a fluid level to build up enough to restrict the inflow from the reservoir, but long enough to keep the well from cycling unnecessarily. Simply put, the idle time should be as long as it can possibly be without hindering production. The assumption is that the optimal idle time is the idle time that still allows the well to achieve the maximum production, with the fewest cycles and pump off strokes possible. Working based off the assumption the goal of the software solution is to optimize the idle time by continually searching for the longest idle time possible that does not hinder the well's production. This gives three variables to monitor, the idle time, production, and cycles for the well. The idle time is the setpoint that tells the well how long to idle. The production is the amount of fluid, typically in barrels per day, that the well is making, and the cycles are the number of times the well goes from on to off back to on. For this algorithm to work, host software must be collecting and storing these thre data points. Using all this data, the software program seeks to optimize the well's downtime to decrease failures and ensure production is not lost.

Description and Application of Equipment and Processes

The solution for autonomous downtime optimization was developed using algorithms that live in a host software program. This program, by design, was made agnostic of the equipment at the wellsite. Utilizing the rod pump controller, the host software takes the inputs of cycles and production and modulates the idle time to reduce the number of cycles without hindering the production. The host software algorithms take the data collected during a twenty-four-hour time period and based on this data make a decision to either increase or decrease the downtime. The software solution will continue to modulate the downtime for the well, if the production begins to decrease the algorithm will begin decreasing the downtime to ensure production is not being hindered by the fluid building up in the casing. Likewise, if the cycles begin to rise the solution will autonomously increase the downtime to avoid high cycles.

The required equipment for the host software solution to optimize the well's downtime is a sucker rod pump system that has a rod pump controller installed with communication either through serial or ethernet connection. The rod pump controller also must report the well's idle time, cycles, and production/runtime. The host software must also have read/write access to the rod pump controller to modulate the setpoints and receive the data for the host software to determine if the setpoint changes were valid. It is also extremely important for the host software to keep record of the historical data for the well so it can be determined over the long term if the downtime changes are truly optimizing the well. Sometimes over short periods of time data can look promising, but unless the solution considers a larger time span, it cannot be certain that the algorithms are accurately optimizing the well.

Once the algorithms were finalized, they were tested against real wells and after a couple of iterations and tweaking the algorithms, they were provided to customers in both the Permian basin and Bakken via software pilots. The results of the pilots were promising and demonstrated huge improvements the well's operation through autonomously optimizing the well's downtime.

Presentation of Data and Results

The initial pilot of the autonomous downtime optimization was run on 100+ wells spread across the Bakken and Permian Basin. The results demonstrated that majority of the wells were able to increase the downtime and reduce daily on-off cycles.

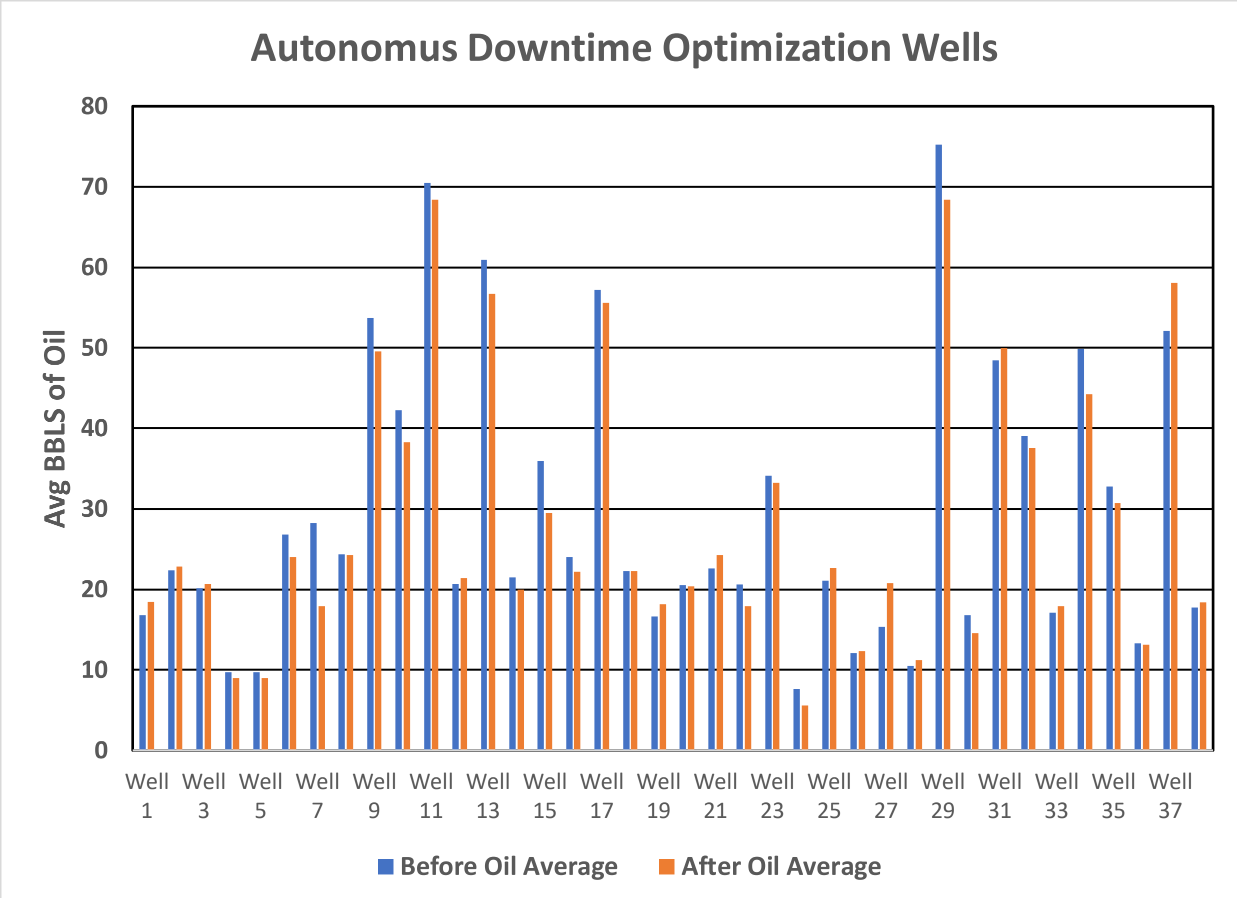

Figure 2: Example of oil production before and after running idle time setpoint optimization.

After removing outliers due to unrelated shutdowns that occurred after the autonomous downtime algorithms were implemented, the average oil production changed from 29.2 bpd to 28.1 bpd. This is a very minimal change in production, however the reduction in cycles was extremely significant averaging a decrease in cycles per day by ~15%. Looking at some examples from the pilots further demonstrates the success from the autonomous downtime optimization.

Case Study 1:

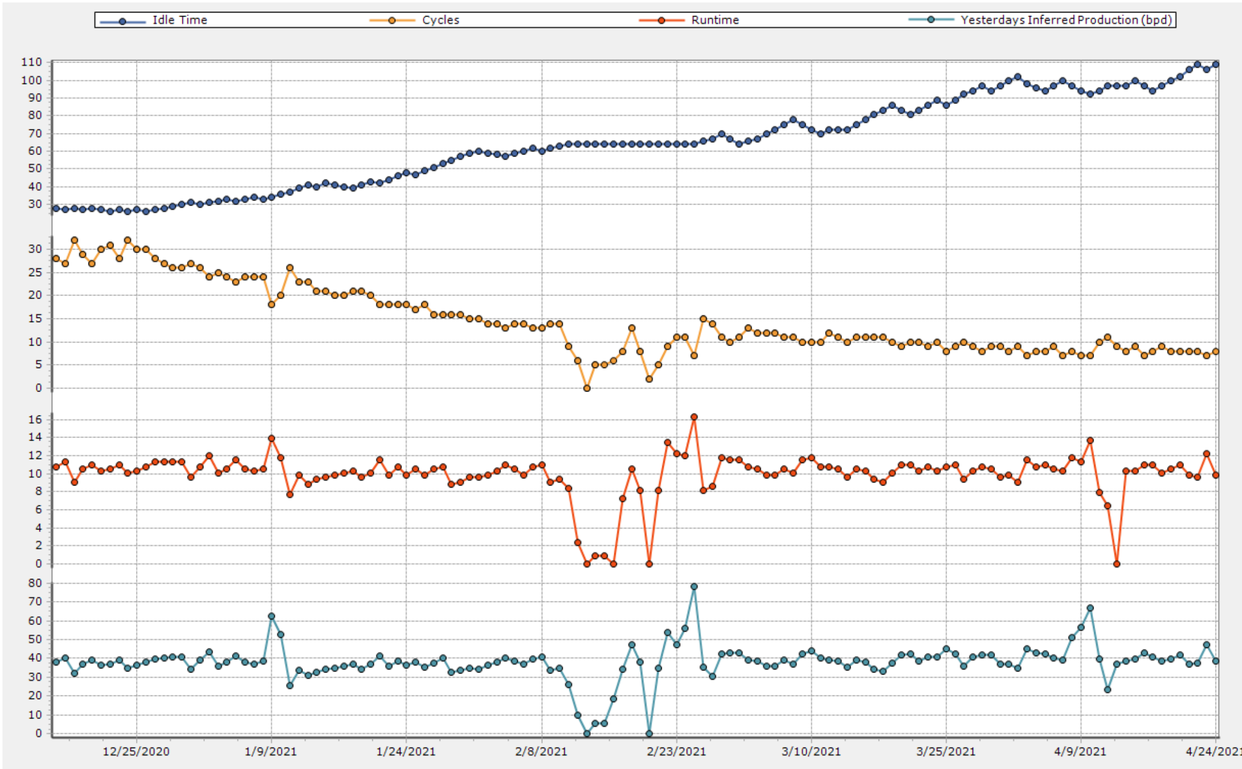

Figure 3: Example historical data from Case Study 1.

Case study 1 is an example of a well experiencing fluid pound in the Bakken. This well began with a downtime of 30 minutes and the autonomous optimization algorithm increased the idle time to 100 minutes (see the top trend). The second trend shows the number of cycles per day and the third and fourth trends show the runtime and inferred production respectively. In this case increasing the downtime had a direct impact on the cycles per day for the well. The cycles dropped from ~30 per day to ~8 per day without having any impact on production. This reduces the number of icomplete fillage pump strokes dramatically and will reduce the failure rate for the downhole equipment, increasing the runtime of the well. The average pump off strokes setting for these wells was 5. Which means this optimization reduced incomplete fillage strokes by over 40,000 strokes in a year. This is a clear example of a well that could have an increase in downtime without hindering production.

Case Study 2:

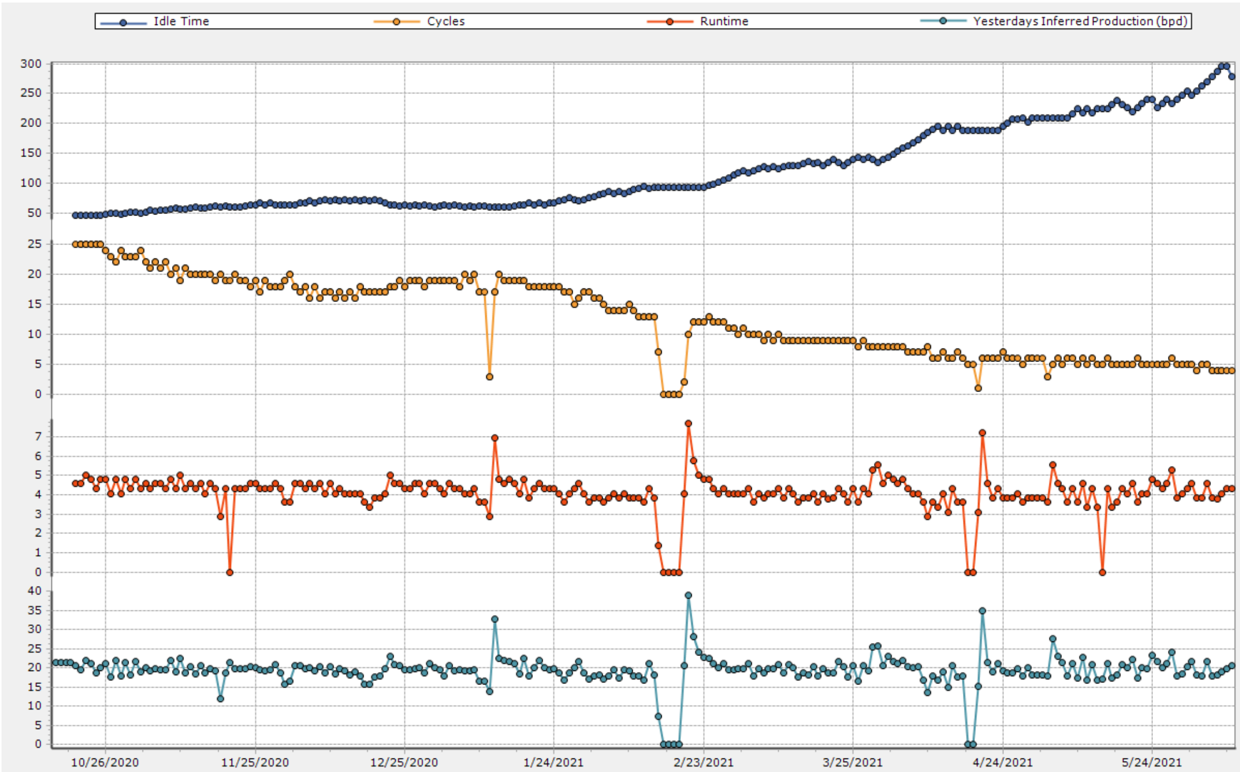

Figure 4: Example historical data from Case Study 2.

Case study 2 is another Bakken well experiencing pump off with fluid pound. This well began with a downtime of 50 minutes. Like Case 1 the autonomous downtime optimization was able to increase the idle time, this time from 50 minutes to 300 minutes. Increasing the idle time brought the cycles per day down to 5 from 25. This is all done without losing production or reducing runtime. This will reduce the number of fluid pound strokes by 36,500 per year (assuming 5 pump off strokes) without hindering the production. By increasing the downtime, the well was still able to produce the same volume of production but run with much fewer incomplete pump fillage strokes which is better for the overall health of the well which will reduce failures.

Case 3:

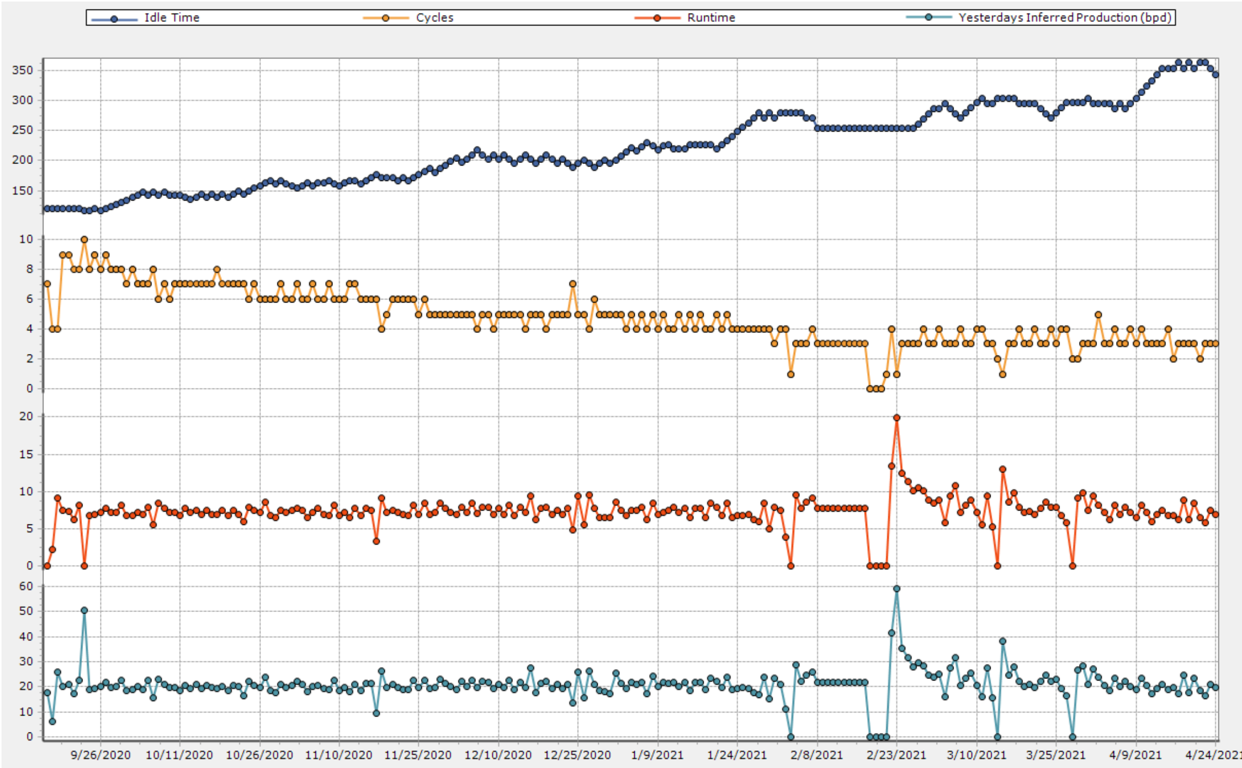

Figure 5: Example historical data from Case Study 3.

Case study 3 is an example of a well in the Permian Basin that was experiencing fluid pound. The idle time was increased from 90 minutes to 345 minutes. Although the idle time was significantly increased, the inferred production remained steady. The downtime optimization algorithms effectively reduced the number of times this well cycled per day, cutting out 20 incomplete fillage strokes per day, without losing any production. This once again achieves the goal of limiting the number of incomplete fillage strokes for the well which will increase the well's run life. This is another case where the autonomous downtime optimization was able to improve the well's performance by increasing the downtime.

Case Study 4:

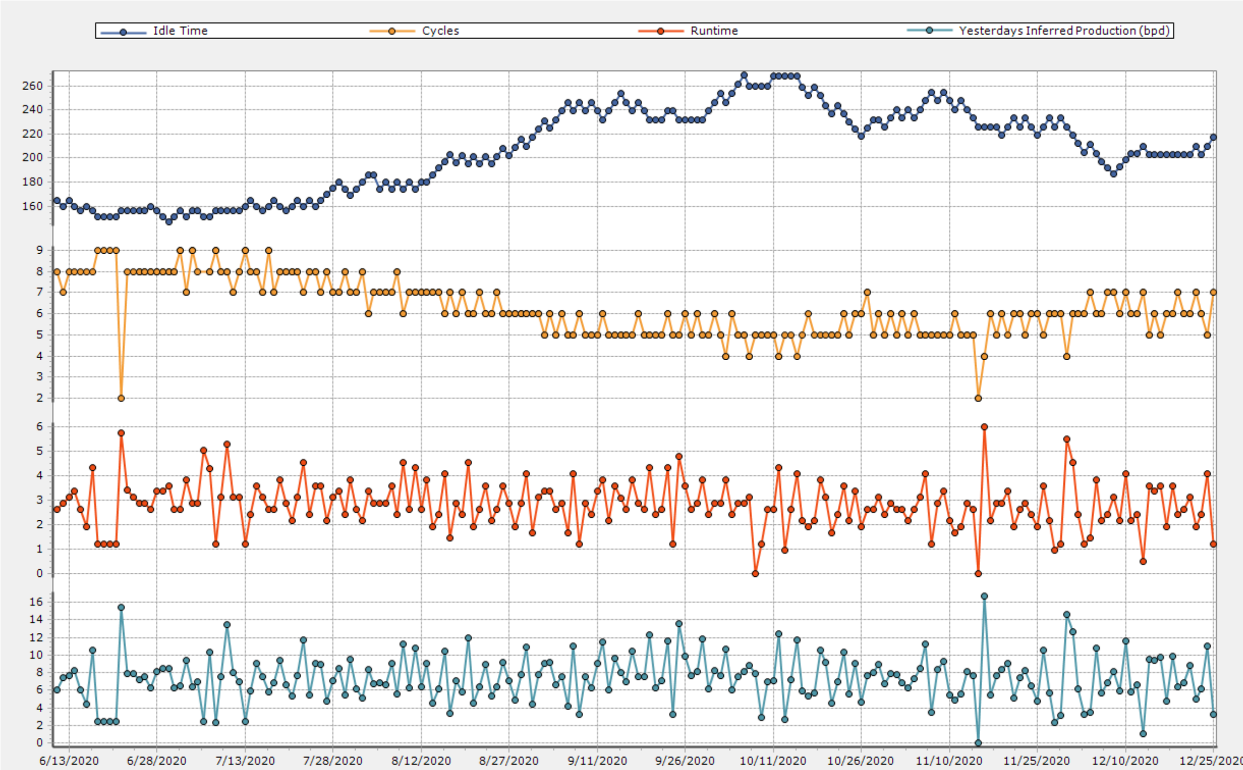

Figure 6: Example historical data from Case Study 4.

Case study 4 is unlike the other case studies. This is an example of a well in the Permian basin with significant gas production leading to gas interference. Optimizing the idle time was not as helpful on this well. The algorithm did increase the well's downtime but started to notice the production was trending down as well. This drove the algorithm to then decrease the downtime. Overall the idle time was increased from 160 minutes to 215 minutes, but this had a much smaller impact on the cycles per day for the well reducing the number of cycles from ~8 to ~5. The production was also not impacted, although the variance in production from day to day is quite large which makes it a little more difficult to analyze, especially when tying to ovserve the impact that idle time has on the production. Although this is not a case where the well was significantly changed or improved by the algorithm, it does validate that the algorithm will not incorrectly change the well's downtime. Not every well can have the idle time optimized, and if it is the case that the current idle time is optimal it is important that the algorithm does not make significant changes to this well. Keeping the downtime at or near where it is for a well that is optimized is as important as optimizing wells that do not have the downtime set at an optimal place.

Conclusion

Optimizing a well's downtime can have major implications on the well's operation. The downtime for a well can be set for a time that is either too short or too long. Having a downtime that is too long will lead to the well losing production from a high fluid level building up in the casing anulus. If the downtime is too short it will cause the well to have unnecessarily high cycles which increases the number of incomplete pump strokes and failure rates. Using a host software solution, it is possible to autonomously optimize the well's downtime and improve the well's operation. In a pilot of over 100 wells it was shown that the downtime optimization had a negligible negative impact on the well's production rate. However, it did have a very significant impact on the well's cycles. On average a well's cycles per day were reduced by ~15%. This is a significant reduction in cycles per day and can aid in reducing failures.

Some specific case studies were explored from this software pilot. In these 4 case studies the downtime optimization algorithm was able to improve operation and reduce cycles without hindering production in 3 of those wells. In the other case the downtime was already set near the optimal downtime and the optimization algorithm made minimal changes to the downtime. This is also seen as a success because the algorithms are not disrupting wells with optimal downtimes. Although it is possible, it was not that case that any of the wells in the initial pilot had downtimes that were too long. This makes sense because many operators are motivated by maximizing production, and in some cases, it appears that this comes at the expense of the well having an increase in pump stokes with incomplete fillaged.

Host software solutions are making huge strides in optimizing a well's downtime. Using legacy rod pump controllers, it is possible to optimize the downtime of a well to reduce failures and increase production. During the pilot it was observed that the software solution was able to increase idle time and reduce cycles in cases where the idle time was set too high. The algorithm also successfully kept the downtime steady in cases where the downtime was optimal. Although there is still plenty of room for growth, host software solutions are making strides in autonomously controlling wells, starting with downtime. Optimizing a wells downtime is a gigantic step in the right direction and a huge win for operators if the algorithms are utilized properly.

Acknowlegement

Thanks to the theta software developers who worked tirelessly to develop these algorithms through the several iterations of generating these algorithms. Carlien, Kevin, Marty, and Chris were all integral to the success of this project. Also, thank you to our customers in the Permian and Bakken that were patient and provided invaluable feedback throughout the testing and development phases of this process.

References

- Bonmmer, P.M. Benefits of Slow Speed Pumping. SPE-103359-JPT
- Kennedy, D.M. and Ghareeb, M. 2017. Ultra-Long-stroke and Intelligent Rod Pumping System for Producing Difficult Wells and/or Fluids. SPE-187028-MS.
- McCafferty, J.F., 1993. Importance of Compression Ratio Calculations in Designing Sucker Rod Pump Installations. SPE -25418-MS.

- Abdalla, R. (2007). *Downhole Problems of Sucker Rod Pumps*. Researchgate.net. https://www.researchgate.net/figure/Downhole-problems-of-sucker-rod-pumps_fig1_332763026.