Master's thesis

NTNU
Norwegian University of Science and Technology
Faculty of Information Technology and Electrical Engineering
Department of Engineering Cybernetics

Torkil Sekse Mollan

Joint Operation of Automatic Managed Pressure Drilling and HeaveLock™

Master's thesis in Cybernetics and Robotics
Supervisor: Ole Morten Aamo
June 2019

NTNU
Norwegian University of
Science and Technology

Torkil Sekse Mollan

Joint Operation of Automatic Managed Pressure Drilling and HeaveLock™

Master's thesis in Cybernetics and Robotics
Supervisor: Ole Morten Aamo
June 2019

Norwegian University of Science and Technology
Faculty of Information Technology and Electrical Engineering
Department of Engineering Cybernetics

Abstract

Offshore drilling technology has come a long way in recent years and is continuously advancing and pushing the limits of what defines an undrillable well. The advancement is fueled by a constant desire of wanting more and to explore even deeper waters. So deep, in fact, that today, we have floating drilling rigs. Floating rigs have introduced new challenges, however, one of which being critical downhole pressure variations caused by the heaving motion of the rig being affected by waves. Managed Pressure Drilling (MPD), a system that is designed to control downhole pressure, has been one of the technologies pushing the limits. However, as is shown in the early part of this thesis, it does not manage to compensate for the heave-induced downhole pressure variations. HeaveLock™, a downhole tool under development by Heavelock AS, aims to eliminate the said problem, enabling drilling even in rough weather conditions. It was their wish to gain knowledge of the joint operation of MPD systems and HeaveLock™ that led to the work presented in this thesis.

Both the MPD control system and the original HeaveLock™ controller was in need of modification before studies of joint operation could be conducted. This thesis shows the development of a proposed modification to allow the MPD system to be implemented on a floating drilling rig, as well as a suggested modification of HeaveLock™'s controller to eliminate drift. A control strategy making use of a moving window algorithm on live measurements is also developed to improve the controller. A drilling simulator is used to study the performance of both modified control systems as well as for joint operation. The implementation of the modified systems in the simulator is also part of this thesis.

The results from this thesis show that, with the proposed modifications, the joint operation of MPD and HeaveLock™ should be feasible.

Sammendrag

Utviklingen av teknologi for boreoperasjoner offshore har kommet en lang vei de siste årene, og fortsetter stadig å skyve grensene på det som definerer at en brønn er umulig å bore. Utviklingen drives av et konstant jag om å ville mere og det å utforske dypere hav. Dette har ført til at vi i dag har flytende borefartøy, som igjen har sørget for nye utfordringer, blant annet kritiske nedihulls trykkvariasjoner som følge av hivbevegelsen til et borefartøy utsatt for bølger. Et system kalt Managed Pressure Drilling (MPD) er utviklet for å regulere nedihulls trykk, og har vært en av de teknologiene som har vært med på å skyve grensene. Men, som denne avhandlingen vil vise, klarer ikke MPD å kompensere for hiv-induserte trykkvariasjoner. HeaveLock™, et nedihullsverktøy under utvikling av selskapet Heavelock AS, skal forsøke å eliminere dette problemet, og dermed muliggjøre boring til og med i tøffe værforhold. Selskapets ønske om å vite mere om hvordan samspillet mellom MPD og HeaveLock™ vil være ga opphav til arbeidet med denne avhandlingen.

Både MPD-systemets styresystem og den opprinnelige regulatoren til HeaveLock™ krevde modifikasjon før det var mulig å studere samspillet mellom systemene. Denne avhandlingen viser utviklingen av en foreslått modifikasjon til MPD-systemet for å muliggjøre implementering på flytende borefartøy, i tillegg til modifikasjon av HeaveLock™ sin regulator for å eliminere drift. En reguleringsstrategi som benytter en såkalt "moving window"-algoritme på sanntidsdata vil også bli utviklet som et forslag til forbedring av HeaveLock™ sin opprinnelige strategi. En boresimulator blir benyttet til å studere ytelsen til de modifiserte systemene i tillegg til samspillet mellom de. Implementasjon av systemene i simulatoren er også en del av denne avhandlingen.

Resultatene fra denne avhandlingen viser at HeaveLock™ og MPD fint kan operere sammen, med lite behov for koordinasjon mellom systemene.

Preface

This thesis is a summary of my work with further development of an existing automatic Managed Pressure Drilling system, designed for strict pressure control on stationary drilling rigs, to enable utilization on floating drilling rigs exposed to harsh weather conditions, as well as examining joint operation of such a system and the downhole tool that a company called Heavelock AS is currently developing. The thesis represents the culmination of my Master's degree in Cybernetics and Robotics at the Norwegian University of Science and Technology (NTNU) in Trondheim and was written from January to June 2019.

The inspiration behind the thesis was the need for modifying a Managed Pressure Drilling system, and implementing it in a drilling simulator, to then perform simulations of joint operation with the downhole tool. This research required an immense amount of lengthy simulations, which was an exhaustive and tiresome, but important process that allowed me to carry out thorough analyses to find answers to the given objectives. The objectives were formulated in cooperation with my supervisor, the co-inventor of Heavelock AS, Professor Ole Morten Aamo.

The simulations presented in this thesis was produced by a simulator I was granted access to by Heavelock AS. All further development, i.e. new strategies and functionality, mentioned in this thesis, I had to implement myself.

I would like to thank my supervisor for his continuing support and reliable guidance during this process. I also wish to thank all involved parties at Heavelock for their constructive input at the weekly meetings, as well as letting me play with their beloved drilling simulator. And to my fellow graduate students: Thank you for being good sparring partners and for always knowing what to do when motivation was low.

Readers are assumed to have prior knowledge of drilling operations and control theory.

Torkil Sekse Mollan	Place and Date

Table of Contents

List of Figures

Abbreviations

BHA Bottom Hole Assembly.

CCS Contiuous Circulation System.

DFT Discrete Fourier Transfrom.

FT Fourier Transfrom.

MD measured depth.

MPD Managed Pressure Drilling.

TVD true vertical depth.

1 Introduction

The oil and gas industry has been drilling wells for years. It is only natural that, with time, the easiest wells are drilled first and the remaining wells become more and more challenging. However, with time, technology is also evolving. The industry is constantly researching new methods and strategies to overcome the challenges that the remaining wells pose. Modern technology is slowly pushing the boundaries, allowing wells that previously have been classified as undrillable to just become more complex problems [1].

Offshore drilling has probably had the greatest technological advances the recent years. Some of the easiest wells have always been those that are easiest to reach, i.e. wells in shallow waters where the drilling rig preferably can stay fixed to the seabed. As technology has evolved, deep waters have become less of a problem. Today, we have floating drilling rigs that can tackle depths that no solid structure ever has before. Floating drilling rigs have introduced new challenges, however, several of which related to handling environmental forces affecting the rig itself, e.g. the heave motion caused by waves. Without a heave compensation system, for example, the drill string would move with the rig and waves, repeatedly impacting the bottom of the wellbore and damaging both well and drilling equipment [2]–[4].

Another technological leap has been the development of systems that enable control of downhole pressure, e.g. manipulation of mud weight, Managed Pressure Drilling (MPD) systems and Continuous Circulation Systems (CCSs). Wells have different degrees of complexity in terms of how narrow their drilling windows are [5]. To prevent dangerous situations, it is important that the downhole pressure is kept within the drilling window at all times. If the pressure falls below the pore pressure, i.e. the lower limit of the drilling window, formation fluids may enter the wellbore. Pressures exceeding the fracture pressure, i.e. the higher limit, may fracture the formation, increasing the chance of drilling fluids exiting the wellbore as well as increasing the formation's permeability [4].

1.1 Background and Motivation

MPD systems are designed to achieve better and more precise control of downhole pressure than just the manipulation of mud weight, allowing drilling of wells with narrow drilling windows [6]. Even though such systems have proved successful when drilling from stationary rigs, research has shown that they have a hard time compensating for heave-induced surge and swab pressures from floating rigs [7], i.e. pressure variations created by the piston-like motion of the drill string within the wellbore [8]. In rough weather, the surge and swab pressures may exceed the drilling window, requiring that drilling must be ceased. This is the inspiration behind HeaveLock™, a new technology currently being developed by Heavelock AS. They aim to eliminate said pressure variations with a downhole tool, enabling drilling from floating drilling rigs even in rough weather conditions [9].

Heavelock's desire is that their tool is designed in such a way that it does not affect the performance of MPD systems, allowing for joint operation of the two technologies without the need of coordination. The main motivation behind this thesis is therefore to study joint operation of automatic MPD systems and HeaveLock™, and to address any conflicts between the two.

1.2 Previous Work

The challenge of finding a solution to the significant pressure variations caused by heave was originally posed by a major oil and gas company. Studies have been performed in the hope of solving this problem for years now, e.g. by trying to utilize topside equipment [8], [10]. Researchers at the Norwegian University of Science and Technology found that it seemed like equipment must be placed downhole to overcome the problem [11], and this revelation led to the founding of Heavelock AS back in 2015.

Heavelock has since then been developing HeaveLock™, a downhole choke valve that aims to compensate for the heave-induced surge and swab pressures by restricting the mud-flow. Their concept is simple; when the downhole pressure increases, restrict the flow and build up pressure in the drill string. When the pressure drops, open the valve and release the built-up pressure. Alongside Heavelock's work with designing HeaveLock™, they have also developed their own drilling simulator (HeaveSim), that they use to produce valuable data for studying the effects of heave when drilling from floating rigs. HeaveLock™ has also been implemented in the simulator such that different control strategies can be tested and verified before deployment.

In a project that was conducted last year, in late 2018, the objective was to gain insight and understanding of MPD and CCSs, and use this knowledge to recommend strategies for implementing HeaveLock™. Both systems were then implemented in HeaveSim to study their performance. It was found that the MPD system being studied was in need of modification to function as expected on a floating drilling rig. The project therefore never came to a conclusion on how to implement HeaveLock™, and is to be seen as preparatory work for this thesis [4].

1.3 Goal and Objectives

The main goal in this thesis is to figure out the effect that joint operation of MPD systems and HeaveLock™ have on each other, and if the systems must be coordinated somehow. However, a previously studied MPD system have shown to be negatively influenced by waves [4], and must be modified to function as if it were to be implemented on a floating drilling rig before we are able to investigate the joint operation of the two systems. The following objectives were therefore formulated in the hope of achieving the goal:

1. Briefly introduce the system that will be studied.
2. Previous work has suggested that the original control strategy of a specific MPD system may amplify heave-induced downhole pressure oscillations, and that MPD is only suited to control slowly varying disturbances and variations as a result of changes in mud-flow [4]. Perform simulations that clarify this problem.
3. Modify the MPD control system in such a way that heave no longer affects its performance. Demonstrate the effect that this modification has on the topside choke and downhole pressure. Previous work suggests that one can achieve a less aggressive choke with positive effects on downhole pressure oscillations. Is this achieved?
4. Utilize HeaveLock™ to reduce downhole pressure oscillations. Does this affect the MPD control system's performance in any way? Can the two systems operate independently of each other, or are they in need of significant coordination?

1.4 Outline

This thesis begins by introducing important theories and concepts that one should know to fully comprehend the problem and goal at hand. The control system and original performance of the specific MPD system that will be modified later are then summarized in Section 3. Simulations and analyses that clarify why the MPD control system is in need of modification are also presented here. The proposed modification is then implemented in Section 4, together with a demonstration of the effects of said modification and the new performance of the modified MPD control system. HeaveLock™ is then introduced in Section 5. Here, HeaveLock™'s original control strategy is presented with examples of how its implemented in the drilling simulator, as well as a performance review. A problem with the original control strategy is also highlighted here, in addition to my proposal for improvement. The study of joint operation is found in Section 6, before the thesis is concluded in Section 7 with proposals for future work.

The appendix consists of three parts, the first of which, Appendix A, being a case example of one of my ideas to improve HeaveLock™'s controller. Appendix B lists a couple of discretization methods that are needed to implement the proposed strategies in the simulator. Appendix C consist of multiple plots showing the performance and responses of the MPD system, both with and without the proposed modifications, and may be of interest for anyone considering implementing said modifications.

2 Theory

2.1 Managed Pressure Drilling

MPD systems are designed to achieve better control over downhole pressure than conventional drilling does. While conventional drilling controls hydrostatic pressure, and thus downhole pressure, by manipulating the density of the drilling fluid (or mud weight), MPD applies back-pressure by manipulating a topside choke valve as well [12], see Fig. 2.1. Manipulation of mud weight is a slow process because of the time it takes to flush the entire mud loop. The response of a MPD system is much faster, where the speed of sound is the limiting factor, see Section 2.2. The back-pressure pump is there to ensure continuous flow through the choke valve, preventing it to saturate and maintaining controllability if the flow were to be interrupted [12], e.g. during connections without CCSs. CCSs are systems that ensure continuous circulation when the physical connection is in progress, eliminating pressure drops created by fading friction loss when circulation is lost and preventing cuttings to settle [13]. When CCSs are used together with MPD systems there is no need for the back-pressure pump to ensure flow through the topside choke [4]. The downhole pressure when using MPD can be found with

$$p_{dh} = p_c + p_{h,a} + p_{f,a}, \tag{2.1}$$

where p_c is the choke pressure, or applied back-pressure, $p_{h,a} = \rho_a g h$ is the annular hydrostatic pressure, $p_{f,a}$ is the annular friction loss, ρ_a is the annular mud density and h is the depth given in true vertical depth (TVD). One distinguishes between drill string and annular hydrostatic pressures and friction losses because the mud weight changes with the amount of cuttings in the annulus, and because the annular walls have different friction factors. MPD systems control downhole pressure by changing the first term in (2.1), the choke pressure, while manipulation of mud weight changes the second term. The friction term changes with mud flow [14], and will be more or less constant with a CCS implemented. The main pump pressure is further found with

$$p_p = p_{dh} + \Delta p_b - p_{h,d} + p_{f,d}, \tag{2.2}$$

where Δp_b is the differential pressure over the bit, $p_{h,d} = \rho_d g h$ is the hydrostatic pressure, $p_{f,d}$ is the friction loss and ρ_d is the mud density within the drill string.

The control system and performance of a specific MPD system have been researched earlier and will be summarized in Section 3.

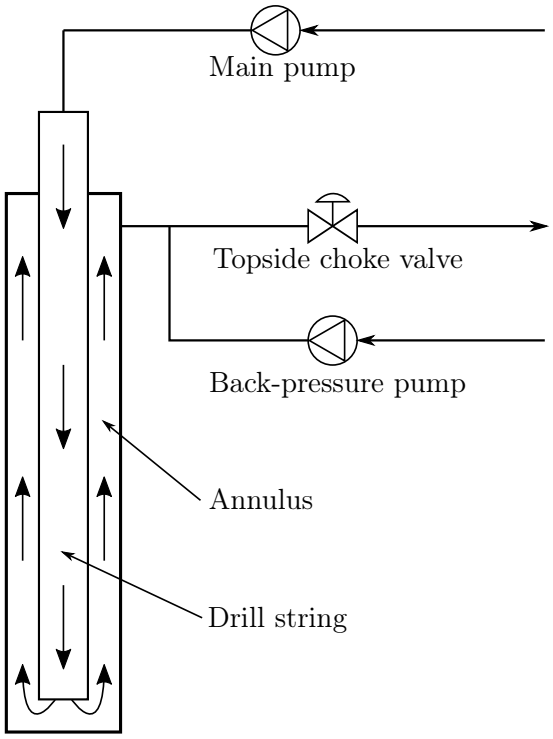

Figure 2.1: MPD overview

2.2 Speed of Sound

Sound waves travel through different mediums at different velocities. Most commonly known is perhaps the speed of sound through air, which is somewhere around 343 m/s, depending on the air temperature and which gases the air is made up of. This means that if you are 343 meters away from a bullet being fired, you will hear it about a second after it was fired. As most modern firearms are able to fire *supersonic* ammunition (bullet travels faster than sound), the bullet probably did not kill you if you heard it was fired. What is also known to most people, is that sound travels faster through denser mediums, i.e. through fluids and solids. In fact, it travels at around 1480 m/s in water, horizontally that is. Generally, the speed increases slightly with pressure and temperature, and thus will change with depth due to the increasing hydrostatic pressure and dropping temperatures [15]. For illustrative purposes, this is assumed negligible.

What is less known, is that sound waves are actually pressure waves, or in other terms, a local compression of a medium propagating through a volume [16]. It is the pressure change in the air around the firearm being fired that produces the sound we can hear. This means that pressures are traveling through a medium at the speed of sound of that medium. Often, when designing or controlling hydraulic systems, we assume that the hydraulic fluid is *incompressible*, i.e. a

pressure change at one end is instantly present at the other. It is usually safe to assume this for small systems, as the distances never come close to making the delay due to the speed of sound in the hydraulic fluid a limiting factor [17]. However, for kilometers deep drilling wells, this delay must be taken into account. The equation for the speed of sound c in a fluid is

$$c = \sqrt{\frac{\beta}{\rho}}, \tag{2.3}$$

where β and ρ is the bulk modulus and density of the fluid, respectively. The time it takes for topside pressure changes to reach the bottom of the well, or similarly, the time it takes for downhole pressure changes to be measured topside is thus

$$T_d = \frac{\text{MD}}{c} = \sqrt{\frac{\rho}{\beta}}\text{MD}, \tag{2.4}$$

where measured depth (MD) is the depth measured along the wellbore. Fig. 2.2 illustrates how downhole pressure waves that travel to the surface at a velocity of c will result in a delayed measurement (or phase lag).

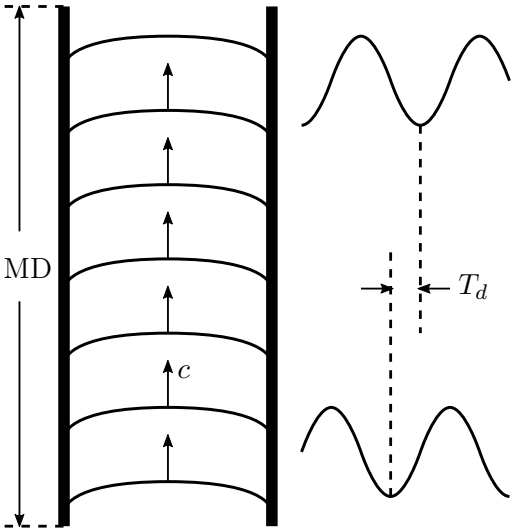

Figure 2.2: Speed of sound

2.3 Effect of Delay

To fully compensate for any downhole pressure changes, one would have to somehow apply the opposite pressure at the same exact moment. If the downhole pressure increases by 2 bar due to changes in downhole conditions, one would have to reduce the pressure by 2 bar at the exact same time to keep the downhole pressure stable. However, because of the limiting speed of sound, there is a delay from the time these changes take place and the time that topside equipment is able to measure said changes and react accordingly. Furthermore, there is another delay before applied pressure changes topside reaches the bottom of the well.

Let us assume that the downhole pressure change due to heave is a perfect sine wave with a period of 15 seconds. To fully compensate for this pressure change, one would have to somehow apply the exact opposite pressure change at the same time. Two sine waves, one with reversed polarity (180 degrees apart), will eliminate one another (the phenomenon is also called destructive interference). However, because of the additional phase lag due to delay, the downhole pressure change may actually get amplified rather than compensated for. Fig. 2.3 shows the phenomenon when a drilling fluid with density of 1946 kg/m^3 and bulk modulus of 21668 bar is used at different MDs. 0 meters has been included in the figure just to show the effect of perfect destructive interference when there is no delay. As the depth increases, it takes time for downhole conditions to reach the top before being measured, and we can see that the optimal applied topside pressure is delayed. The applied topside pressure is further delayed before it arrives downhole. At 3957 meters there is perfect constructive interference, i.e. both the uncontrolled downhole pressure and the arriving applied topside pressure are in phase and the result is a wave with maximum amplification. We can see that already around a depth of 1500 meters, the resulting downhole pressure has a bigger amplitude than 2 bar, meaning the downhole pressure is being amplified. This is, of course, a very simplified simulation, just to illustrate the effect of the delay. In reality, there are probably several other variables that will affect the results. The simulation does not, for example, take into account that the increased resulting controlled downhole pressure will increase the applied topside pressure as well, only to increase the resulting controlled downhole pressure even more.

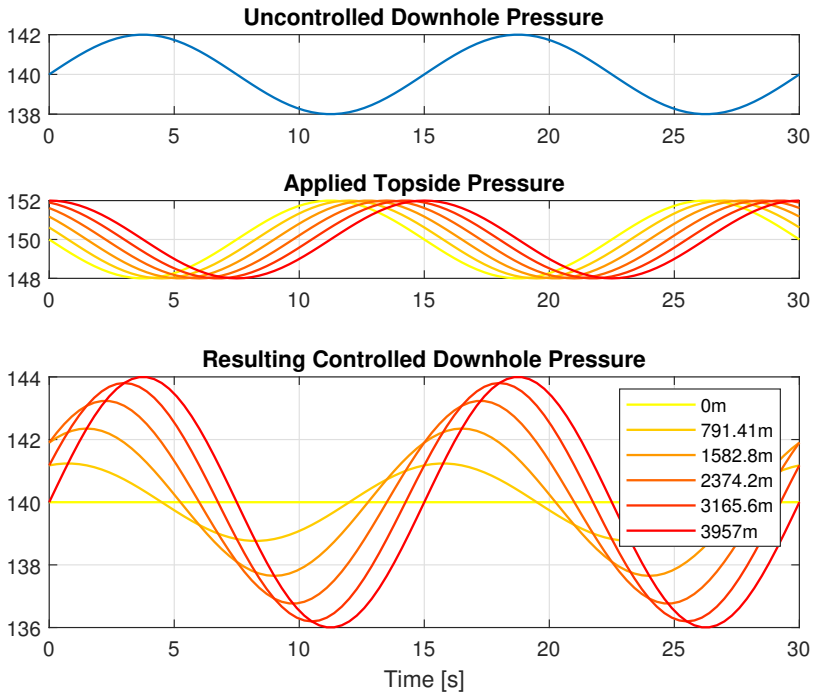

Figure 2.3: Interference at different measured depths (MDs)

2.4 Filters

2.4.1 Low-pass Filters

A low-pass filter is simply a filter that will allow low frequencies to pass and reduce the amplitude of the higher ones. This fits the purpose of the MPD system summarized in Section 3 quite well, because it is designed to control slowly varying conditions (low-frequency disturbances), but should not try to control the oscillations that come from heave (higher frequency disturbances) [4]. A simple first order low-pass filter can be written in Laplace form as

$$h_{lp}(s) = \frac{output}{input} = \frac{K}{\tau s + 1} = \frac{K\omega_c}{s + \omega_c}, \tag{2.5}$$

where K is the gain for the frequencies that are allowed to pass, and τ is the filter's time constant. The time constant or cutoff frequency $\omega_c = 1/\tau$ is what defines the separation between passed and reduced frequencies, i.e. all lower frequencies than the cutoff frequency will pass, while the higher ones will be reduced.

A second order low-pass filter has the same objective as a low-pass filter of first order, but its performance is better, i.e. it reduces the higher frequencies even more, creating a better separation between passed and reduced frequencies. A second order filter has slower responses

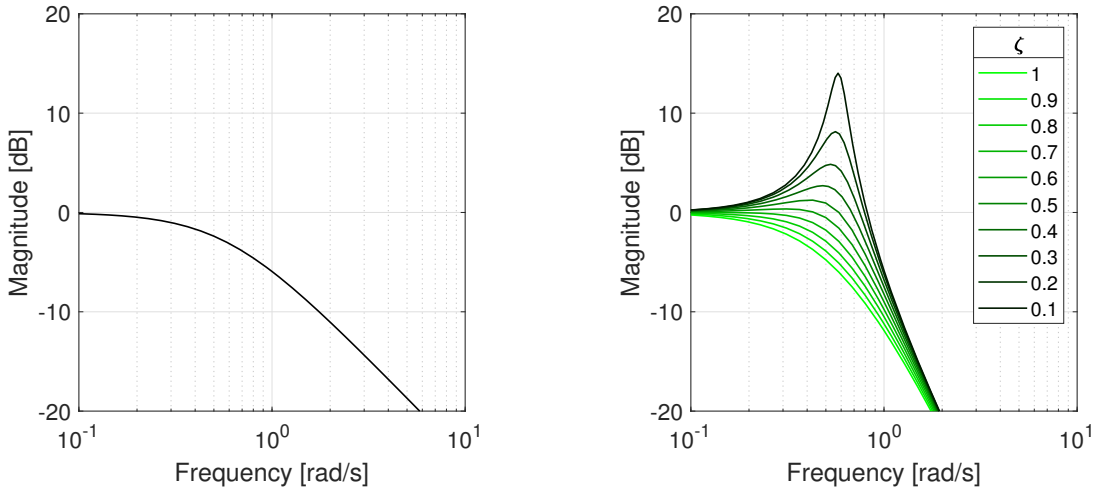

(a) First order (b) Second order

Figure 2.4: Order comparison and cutoff frequency analysis of low-pass filters

than one of first order, however, i.e. it adds additional phase lag. It is typically written as

$$h_{lp}(s) = \frac{K\omega_c^2}{s^2 + 2\zeta\omega_c s + \omega_c^2}, \tag{2.6}$$

where ζ is a damping ratio, describing the filter's behaviour around the cutoff frequency. The frequency response of a first and second order low-pass filter with different damping ratios is plotted in Fig. 2.4. All filters had a gain equal to 1 and a cutoff frequency of 0.58 rad/s. Low damping ratios result in amplification of frequencies around the cutoff frequency and is not desirable in most filter applications. The damping ratio should be experimented with in cases where the frequencies we wish to attenuate is close to the ones we need. The figure also shows that the second order filter has a steeper curve for higher frequencies. Even steeper curves can be achieved with higher-order filters at the expense of additional phase lag.

2.4.2 High-pass Filters

A high-pass filter is exactly the opposite of a low-pass filter, i.e. it allows high frequencies to pass, but attenuate lower frequencies. Such a filter can be used in applications where only the higher frequency disturbances are of interest and not slow disturbances such as drift. High-pass filters can also be of different orders, where first and second order filters are written as

$$h_{hp}(s) = \frac{Ks}{s + \omega_c} \qquad \text{and} \qquad h_{hp}(s) = \frac{Ks^2}{s^2 + 2\zeta\omega_c s + \omega_c^2}, \tag{2.7}$$

respectively. The frequency response of high-pass filters are symmetrically equal to the low-pass filter frequency response, i.e. about the line where $\omega = \omega_c$, see Fig. 2.5.

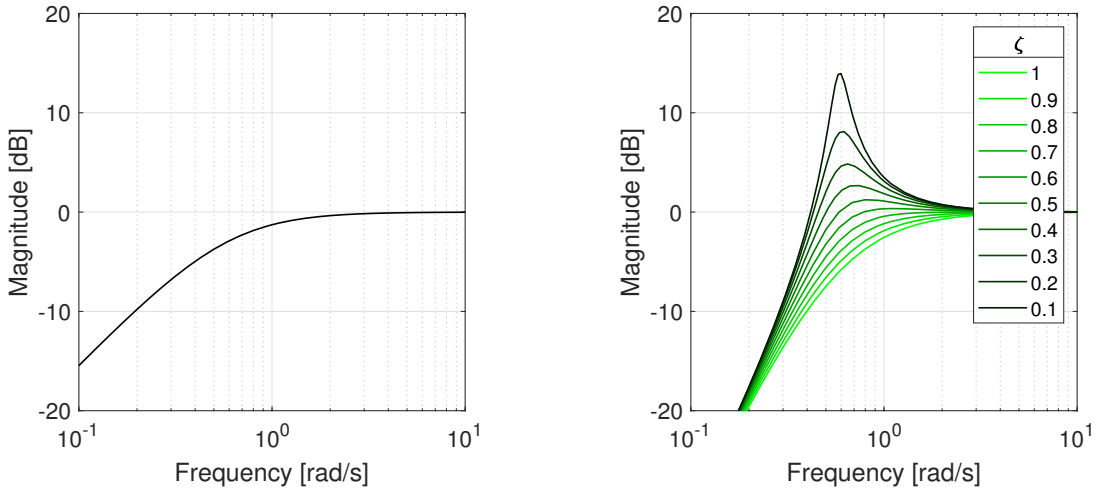

(a) First order (b) Second order

Figure 2.5: Order comparison and cutoff frequency analysis of high-pass filters

2.4.3 Band-pass Filters

Band-pass filters can be used if we are only interested in a specific range of frequencies, i.e. a passband. Some band-pass filters are just cascaded high- and low-pass filters where the passband is defined by the filters respective cutoff frequencies. A simple second-order band-pass like this can be defined from the first order high-pass filter in (2.7) and (2.5), i.e.

$$h_{bp}(s) = K\frac{s}{s+\omega_{c1}} \cdot \frac{\omega_{c2}}{s+\omega_{c2}}, \tag{2.8}$$

where ω_{c1} and ω_{c2} are the lower and higher cutoff frequencies of the passband, respectively, and the bandwidth is approximately $w_{bw} = w_{c2} - w_{c1}$. Note that this filter will attenuate the frequencies within the passband slightly if the passband is narrow. If this is the case the following filter may be a more suitable solution, i.e.

$$h_{bp}(s) = \frac{\omega_{bw}s}{s^2 + \omega_{bw}s + \omega_c^2}, \tag{2.9}$$

where ω_{bw} and ω_c is the bandwidth and center frequency of the passband, respectively. Fig. 2.6 shows a comparison of the two types of band-pass filters. Note how the cascaded high- and low-pass filter in Fig. 2.6a attenuates frequencies inside the passband slightly if the bandwidth

is narrow, while the bandwidth-defined filter in Fig. 2.6b keeps the center frequency at zero gain.

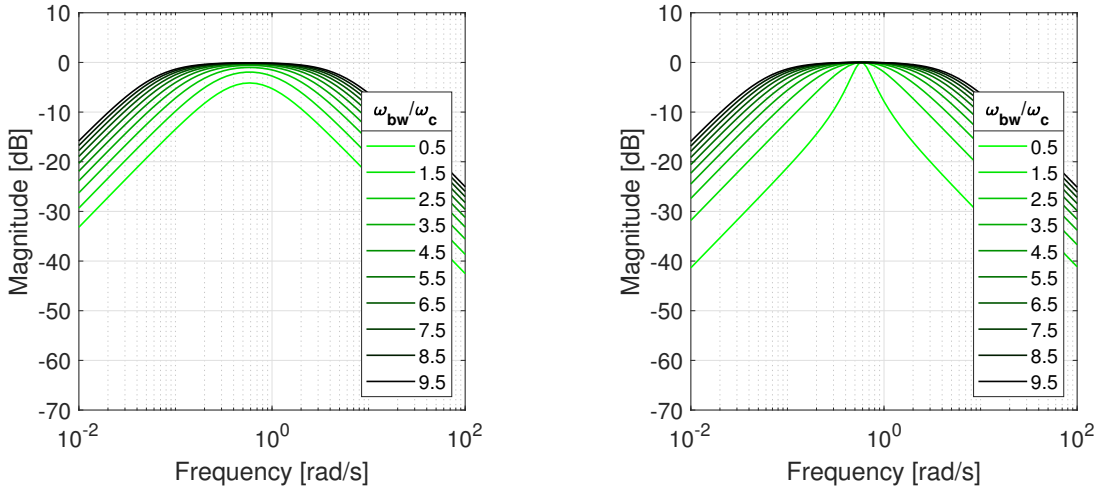

(a) Cascaded high- and low-pass

(b) Bandwidth-defined

Figure 2.6: Bandwidth analysis of band-pass filters

If we wish to attenuate the frequencies outside of the passband even more, a higher-order Butterworth band-pass filter can be a great option, e.g. one of fourth order, i.e. [18]

$$h_{bp}(s) = \frac{\omega_{bw}^2 s^2}{s^4 + \sqrt{2}\omega_{bw}s^3 + (2\omega_c^2 + \omega_{bw}^2)s^2 + \sqrt{2}\omega_{bw}\omega_c^2 s + \omega_c^4}. \tag{2.10}$$

2.4.4 Band-stop/Notch Filters

For some applications, the unwanted range of frequencies (or frequency spectrum) may be within the controller's frequency domain, i.e. the controller needs control over both lower and higher frequencies than those described by the frequency spectrum. This could, for instance, be applications that have to react to sudden changes very quickly, but at the same time the controller is heavily influenced by some disturbance that does not necessarily affect the performance by much, but could cause major wear and tear on actuators, requiring more frequent maintenance, and resulting in larger expenses. In such applications, a band-stop/notch filter could be a good idea to implement. A notch filter is designed to reduce a specific band of frequencies, allowing both the lower and higher frequencies of the limiting stopband to pass. Such a filter can be used to reduce only specific unwanted frequencies and is typically used to dampen the humming sound from the power grids influencing measurements. A general notch filter can be written on

the form

$$h_n(s) = \frac{s^2 + 2\zeta\omega_{bw}s + \omega_c^2}{s^2 + 2\omega_{bw}s + \omega_c^2}, \tag{2.11}$$

where ζ defines the magnitude of the notch (see Fig. 2.7a), and ω_c and ω_{bw} is the center frequency and bandwidth of the stopband, respectively. The bandwidth defines the width of the stopband, as shown in Fig. 2.7b. Note that the different bandwidths in the figure are shown as a ratio with respect to the cutoff frequency. The notch filters in the figure all have a center frequency of 0.58 rad/s.

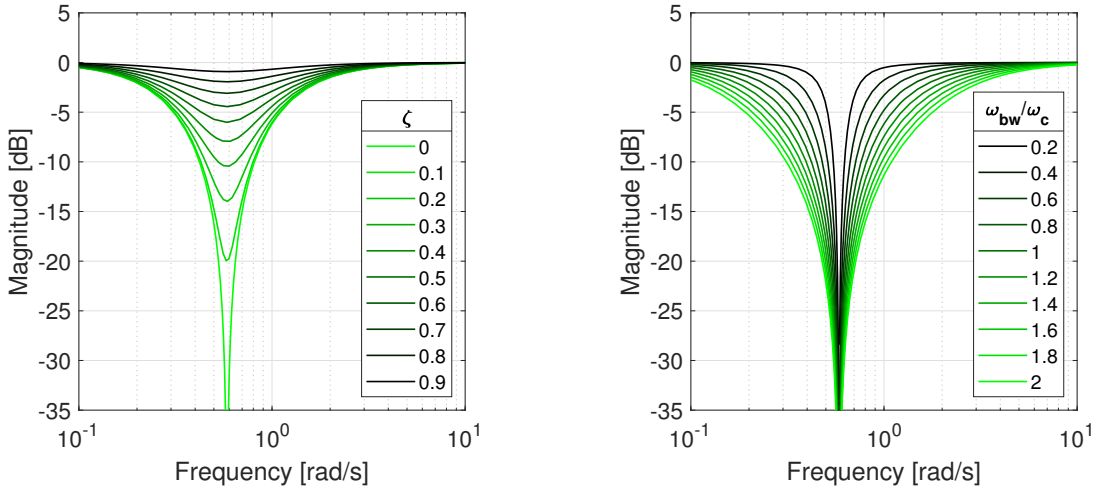

(a) Damping ratio

(b) Bandwidth

Figure 2.7: Parameter analysis of notch filters

2.4.5 Cascaded Low-pass and Notch Filters

Although a notch filter can be a solution for some applications, there could still be high frequencies that either may affect the controller's performance or cause wear and tear, e.g. noisy measurements. It would then be a good idea to use cascaded low-pass and notch filters, i.e.

$$h_f(s) = h_{lp}(s)h_n(s), \tag{2.12}$$

where $h_{lp}(s)$ and $h_n(s)$ can be defined from (2.5) and (2.11). The cutoff frequency of the low-pass filter must be chosen higher than the center frequency of the notch filter. Fig. 2.8 shows the frequency response of a cascaded first order low-pass and notch filter with different cutoff frequencies for the low-pass filter. The notch filter used in the figure has a center frequency of 0.58 rad/s. The figure shows that the low-pass filter will dominate the notch filter as the cutoff

frequency decreases. If the cutoff frequency is less than the center frequency, the notch filter is most probably not necessary and just a simple low-pass filter will work just as well.

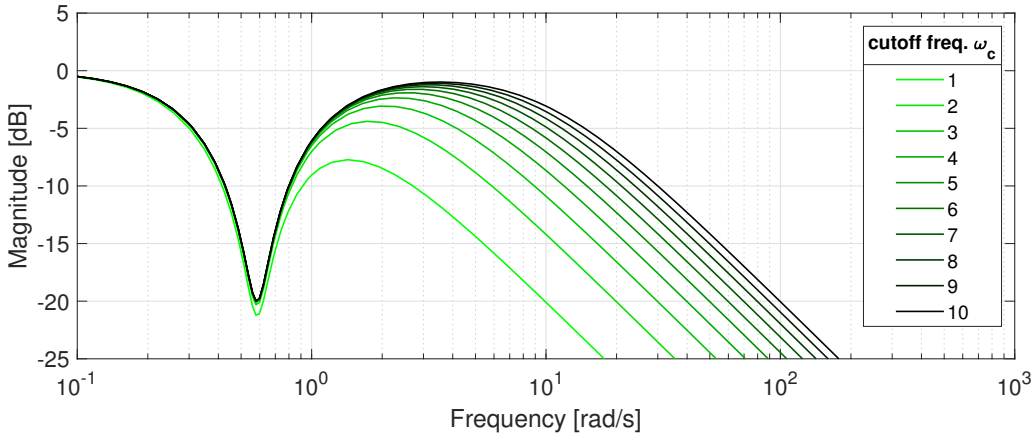

Figure 2.8: Cutoff frequency analysis of cascaded low-pass and notch filter

Some times the frequency spectrum that we wish to attenuate is wider than expected. In such cases, a single notch filter may not always be enough to cover the whole spectrum. If the frequency range of the spectrum exceeds the width of a single notch filter, multiple notch filters in cascade may be a suitable solution, i.e.

$$h_f(s) = \prod_{i=1}^{m} h_{n,i}(s), \qquad \text{where} \qquad h_{n,i}(s) = \frac{s^2 + 2\zeta_i\omega_{bw,i}s + \omega_{c,i}^2}{s^2 + 2\omega_{bw,i}s + \omega_{c,i}^2}, \tag{2.13}$$

and ζ_i, $\omega_{bw,i}$ and $\omega_{c,i}$ can be chosen individually for each notch filter to fit your needs [19]. Even though a single filter may be sufficient in simulations where the frequency spectrum is known, using multiple notch filters should be considered for real implementations where the frequency spectrum is uncertain.

An example of a cascaded low-pass and notch filter that may be suitable for our purposes is a filter suggested by Micheal J. Grimble and Micheal A. Johnson in 1989 [19], [20], to remove typical frequencies found in waves. The filter consists of a low-pass filter in cascade with three unique notch filters and is defined as

$$h_f(s) = \frac{1}{\tau s + 1} \cdot \prod_{i=1}^{3} \frac{s^2 + 2\zeta_i\omega_{bw,i}s + \omega_{c,i}^2}{s^2 + 2\omega_{bw,i}s + \omega_{c,i}^2}, \tag{2.14}$$

where the respective center frequencies $\omega_{c,i}$ and bandwidths $w_{bw,i}$ are identical and equal to 0.4, 0.63 and 1.0 rad/s, respectively. The time constant is chosen as $\tau = 0.1$ s.

2.4.6 Discrete Filters

All filters described above are defined as transfer functions in the Laplace domain and must be discretized before they can be implemented in modern digital applications.

First Order Low-pass Filter

A first order low-pass filter is written as

$$h_f(s) = \frac{y}{u}(s) = \frac{K}{\tau s + 1}, \tag{2.15}$$

where y and u is the filter's output and input, respectively. To discretize simple first order transfer functions, the first step is usually to translate them to differential equations with use of inverse Laplace, i.e.

$$\tau \dot{y} + y = Ku. \tag{2.16}$$

Multiple numerical methods can be used to discretize differential equations. The most commonly used method for discretizing basic filters is probably backward Euler (B.2), i.e.

$$\tau \dot{y} + y = Ku \tag{2.17a}$$

$$\tau \cdot \frac{y_k - y_{k-1}}{h} + y_k = Ku_k \tag{2.17b}$$

$$y_k = \frac{\tau}{\tau + h} y_{k-1} + \frac{h}{\tau + h} Ku_k \tag{2.17c}$$

$$y_k = (1 - \alpha) y_{k-1} + \alpha K u_k, \tag{2.17d}$$

where u_k and y_k is the filter's input and output for time step k, respectively, and $\alpha = \frac{h}{\tau + h}$. Eq. (2.17d) is probably the easiest discrete low-pass filter to implement, because it requires the least amount of code.

Higher Order Filters

Discretizing filters of higher order the same way with backward Euler can sometimes be troublesome, as solving for x_k can be tricky and time-consuming. However, if we convert the transfer function to state-space form, i.e.

$$\dot{\boldsymbol{x}} = \boldsymbol{A}\boldsymbol{x} + \boldsymbol{b}u \tag{2.18a}$$

$$y = \boldsymbol{c}^\top \boldsymbol{x} + du \tag{2.18b}$$

we can use forward Euler (B.1) to discretize the system as shown in Appendix B.2. We then have

$$\boldsymbol{x}_{k+1} = (\boldsymbol{I} + h\boldsymbol{A})\boldsymbol{x}_k + h\boldsymbol{b}u_k, \tag{2.19}$$

where we just reduce the time step by one to end up with

$$\boldsymbol{x}_k = (\boldsymbol{I} + h\boldsymbol{A})\boldsymbol{x}_{k-1} + h\boldsymbol{b}u_{k-1} \tag{2.20a}$$

$$y_k = \boldsymbol{c}^\top \boldsymbol{x}_k + du_k. \tag{2.20b}$$

An easy way to convert an arbitrary proper transfer function

$$h(s) = \frac{b_n s^n + \ldots + b_1 s + b_0}{s^n + \ldots + a_1 s + a_0}, \tag{2.21}$$

of order n to the state-space form in (2.18) is to use one of many canonical forms, e.g. the observable canonical form where [18]

$$\dot{\boldsymbol{x}} = \begin{bmatrix} -a_{n-1} & 1 & 0 & \ldots & 0 \\ -a_{n-2} & 0 & 1 & \ldots & 0 \\ \vdots & \vdots & \vdots & & \vdots \\ -a_0 & 0 & 0 & \ldots & 0 \end{bmatrix} \boldsymbol{x} + \begin{bmatrix} b_{n-1} - a_{n-1}b_n \\ b_{n-2} - a_{n-2}b_n \\ \vdots \\ b_0 - a_0 b_n \end{bmatrix} u \tag{2.22a}$$

$$y = \begin{bmatrix} 1 & 0 & 0 & \ldots & 0 \end{bmatrix} \boldsymbol{x} + b_n u. \tag{2.22b}$$

3 MPD Control System

As mentioned in Section 2.1, MPD systems aim to control downhole pressure by manipulation of a topside choke valve, together with a back-pressure pump, depending on if a CCS system is implemented or not. In some of my previous work with a specialization project, I investigated the control system in a MPD system developed by researchers at Equinor, as well as implemented it in HeaveSim. This system will only be summarized here, and interested readers are referred to my project report or the original paper for further details [4], [14].

The MPD system consist of two main subsystems, i.e. a downhole estimator and a choke controller. The estimator is used to find estimates of, among other states, the annular hydrostatic pressure and friction loss in (2.1), such that a desired choke pressure can be found with

$$p_c^d = p_{dh}^d - \hat{p}_{h,a} - \hat{p}_{f,a}, \tag{3.1}$$

where p_{dh}^d is the desired downhole pressure found from a pre-defined *drilling window*, i.e. a window defined by minimum and maximum pressure limits as a function of TVD to ensure safe drilling and avoid dangerous situations at all depths. The desired choke flow is then calculated as

$$q_c^d = \hat{q}_b + \hat{q}_f + q_{bpp} - A_d v_d + \frac{V_a}{\beta_a}\left(k_p(p_c - p_c^d) - \dot{p}_c^d\right), \tag{3.2}$$

where $\hat{q}_b$ is the estimated bit flow, $\hat{q}_f$ is the estimated formation flow, both found from the downhole estimator, q_{bpp} is the back-pressure pump flow, A_d is the cross-sectional area of the drill string, v_d is the velocity of the drill string with respect to the wellbore wall, V_a is the annular volume, β_a is the bulk modulus of the mud in the annulus, k_p is a control parameter and p_c is the measured choke inlet pressure. The derivative of the desired choke pressure $\dot{p}_c^d$ is assumed known, either from a reference model or by differentiating (3.1) directly. Reference models are designed to generate smooth reference trajectories from sudden setpoint changes and contains information about the derivative of the trajectory. Once the desired choke flow is found, the choke valve control input can be found with

$$u_c = A^{-1}\left(\frac{q_c^d}{C_c\sqrt{p_c - p_{c0}}}\right), \tag{3.3}$$

where $A(\cdot)$ is the characteristic function of the valve, C_c is the valve constant, and p_c and p_{c0} are the pressures at the inlet and outlet of the choke valve, respectively.

3.1 Performance

The MPD system summarized above was implemented in a drilling simulator during the work with my specialization project. Multiple simulations were run to study the performance of the MPD system when drilling from both stationary and floating drilling rigs, many of which are found in Appendix C. The simulations showed that the MPD system did not perform well on floaters, as it tried to compensate for any heave-induced pressure oscillations downhole without any luck. Given that HeaveLock™ is to be implemented on floaters, implementing a filter to attenuate these oscillations, in such a way that the MPD system does not act on them, was suggested as future work and led to parts of this thesis [4].

To have a basis for comparing later implementations with, we need to perform some simulations showing the performance of the MPD system as it is. We need simulations of how it performs both with and without heave, as well as with and without CCSs. Simulations with and without heave can also be a measure for how well the system performs on floating versus stationary drilling rigs. The simulations should show responses from varying conditions like main pump rates and setpoint changes as this is what the MPD system is designed for. We also need simulations of connections in both calm and rough conditions, as this is what is relevant for the soon to be designed filter and of course for HeaveLock™.

An analysis of the MPD system's performance is summarized in the following sections. Simulations without CCS have shown that having a constant mud-flow from the back-pressure pump gives the best results. All simulations without CCS implemented is therefore done with a constant back-pressure pump flow. The main pump flow has also been reduced in these cases to have a similar total choke flow as with CCS. The back-pressure pump was off when CCS was implemented. In the simulations of connections, the drill string has been hoisted up 5 meters for 10 seconds from $t = 25$ s and lowered back down again for 10 seconds from $t = 390$ s. The main pump ramps down at around $t = 40$ s and ramps back up again at around $t = 300$ s.

All simulations are done of a vertical well to minimize any uncertainties related to stiction, i.e. the drill string should be hanging freely within the wellbore. The TVD was around 4000 meters in all cases. The data used to produce the simulation results represents typical data for the Norwegian continental shelf.

3.1.1 Pump Flow Ramp Response

The response of a ramp in the main pump flow of the original MPD system is shown in Fig. 3.1, both with and without CCS implemented. The main pump ramps down 30% at $t = 100$ s and back up again at $t = 300$ s in both cases, as shown in Fig. 3.1c. From Fig. 3.1a we see that there is a slight change in downhole pressure when ramping is evident. The MPD system seems to have no problems with maintaining the desired downhole pressure at lower pump rates. As mud flow through the choke valve decreases, so does the differential pressure over the valve. By decreasing the choke's opening, the differential pressure, and thus downhole pressure, is maintained. The non-smooth choke opening seen in Fig. 3.1d is most probably due to the MPD system trying to compensate for the noisy main pump flow measurements.

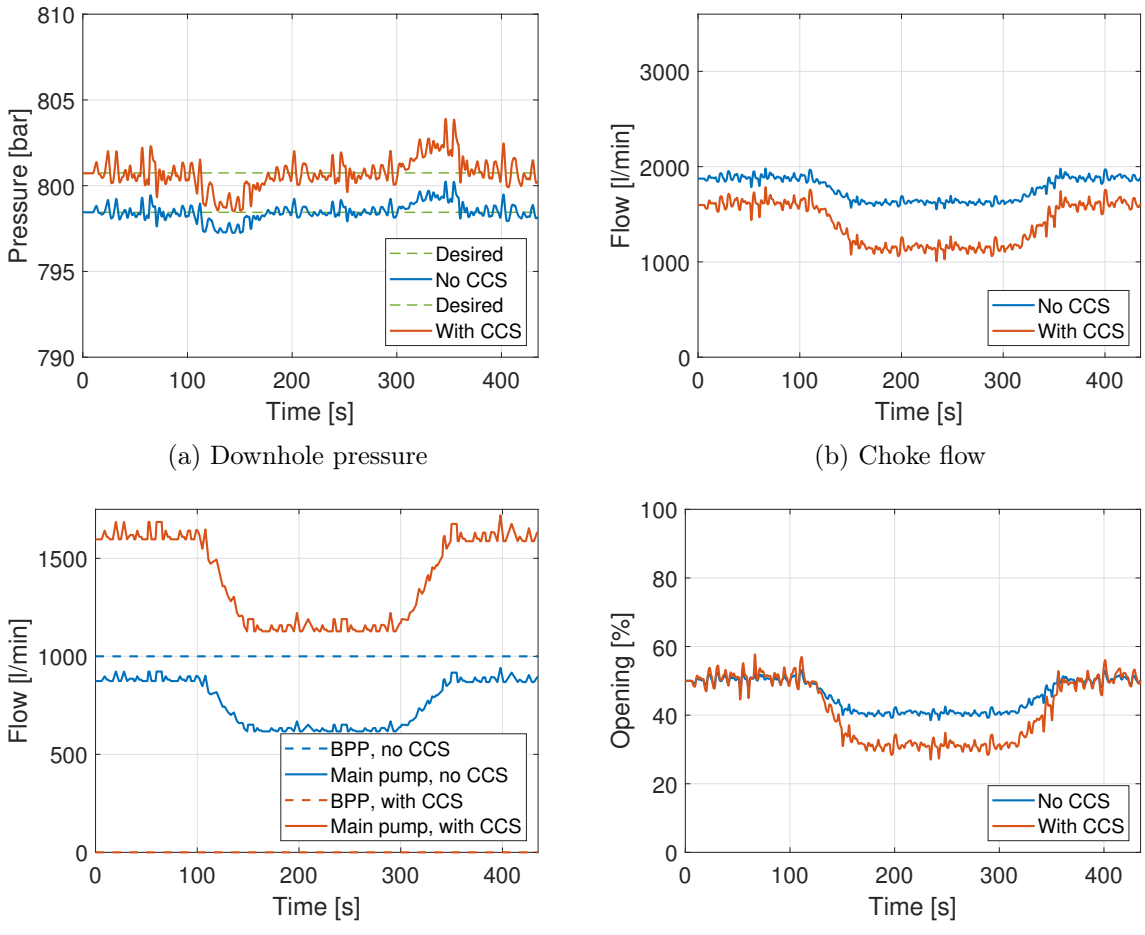

(a) Downhole pressure

(b) Choke flow

(c) Pump flows

(d) Choke opening

Figure 3.1: Pump flow ramp response of original MPD system

3.1.2 Setpoint Step Response

In Fig. 3.2 we see the reference step response of the original MPD system, and again, both with and without CCS implemented. Fig. 3.2a shows the response of a positive step of 10 bar in desired downhole pressure at $t = 150$ s and a step back down at $t = 250$ s. The only noticeable difference with and without CCS implemented seems to be the more noise-exposed choke valve. Again, this is probably because the main pump flow with CCS implemented is somewhat more noisy, due to the fact that it is just an upscaled version of the main pump flow without CCS. However, the response is quick, with no noticeable overshoot in downhole pressure. The choke valve quickly closes to increase the downhole pressure, temporarily decreasing the choke flow, before stabilizing on a stationary smaller opening. The opposite happens when the desired downhole pressure decreases.

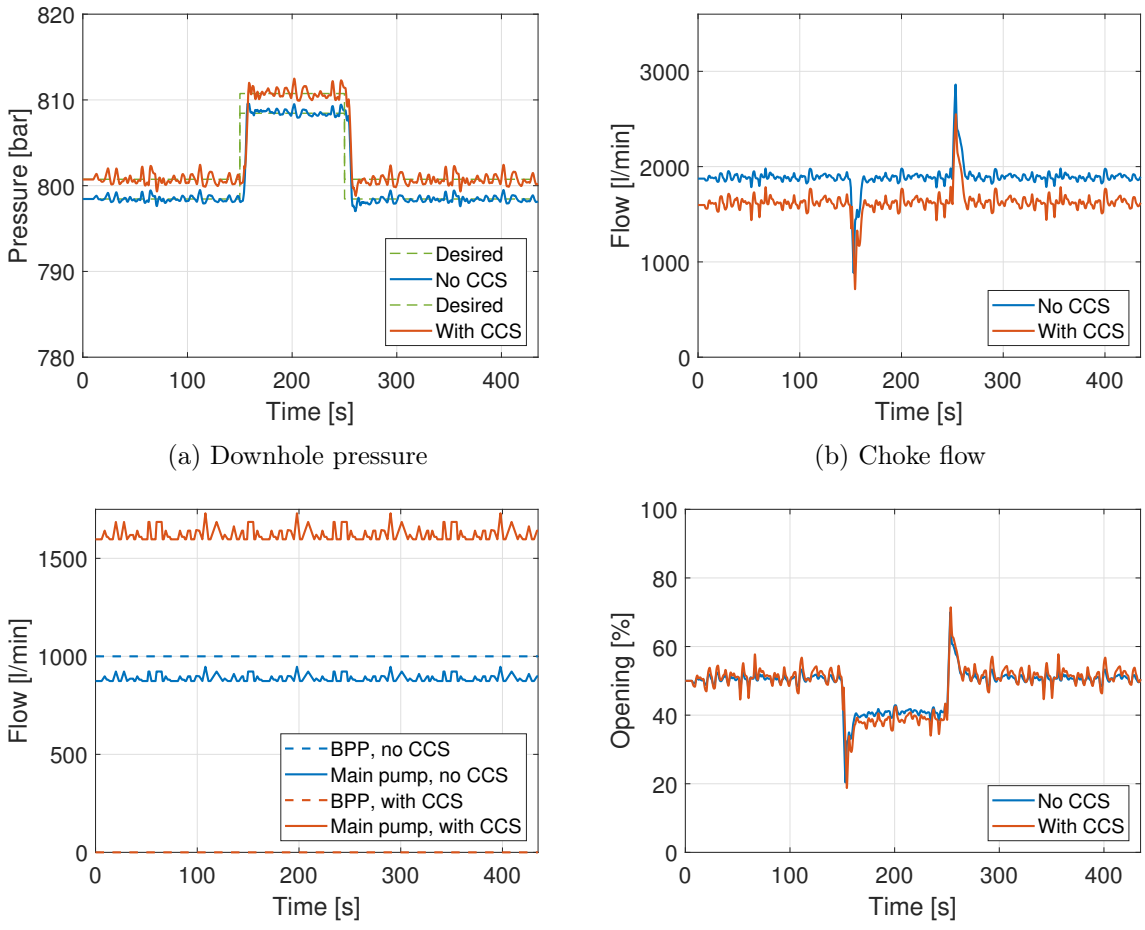

(a) Downhole pressure

(b) Choke flow

(c) Pump flows

(d) Choke opening

Figure 3.2: Setpoint step response of original MPD system

3.1.3 Connection Response in Calm Conditions

Fig. 3.3 shows the response when a connection is done with MPD in calm conditions, both with and without CCS. The change that CCS introduces is seen in Fig. 3.3c, where CCS has a constant main pump flow and no back-pressure flow. Without CCS the main pump must be turned off to perform the connection, and as we can see from Fig. 3.3a, this affects the downhole pressure quite a bit. The choke valve opening (Fig. 3.3d) has to decrease its opening to 20% to compensate for the reduced flow (Fig. 3.3b). In comparison, the response with CCS still operates around the same value, i.e. around 50% choke opening. In both cases the MPD system manages to keep the downhole pressure at its desired value.

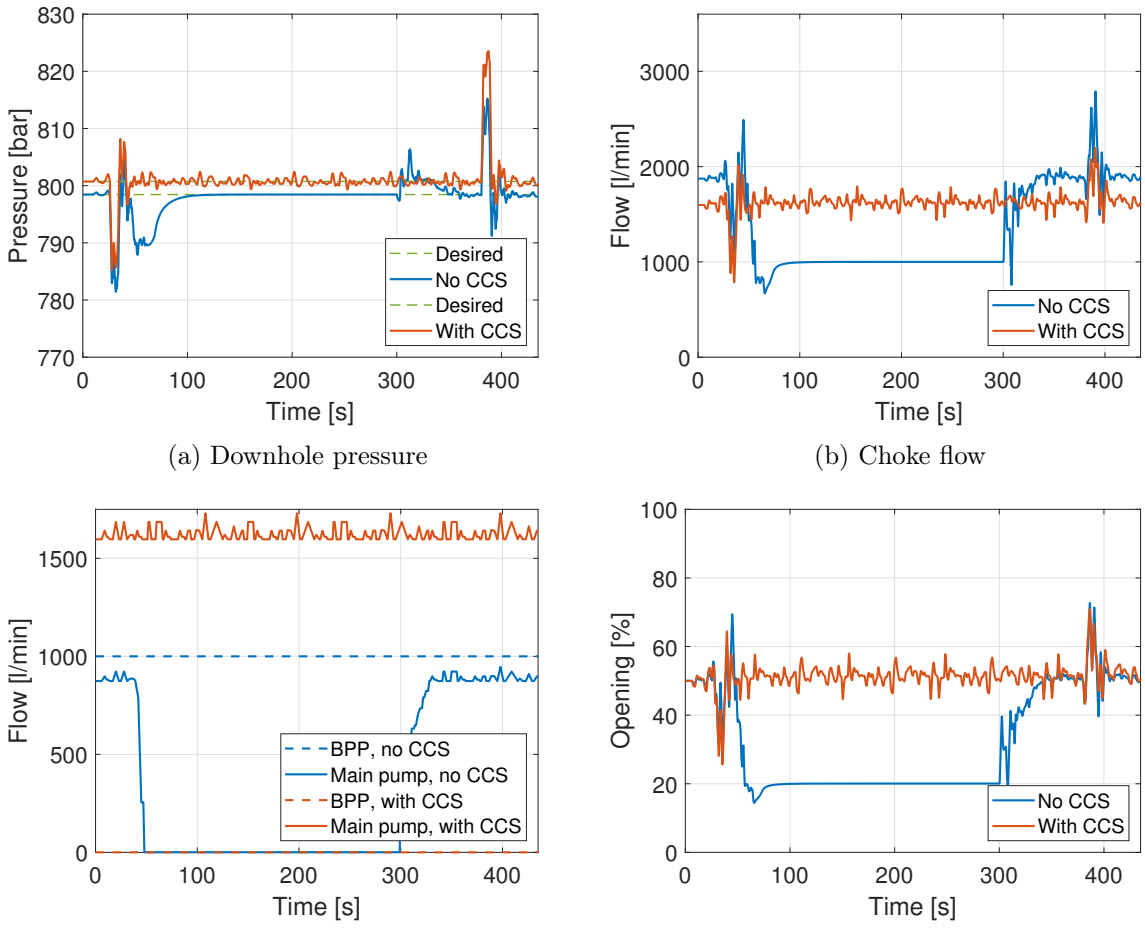

(a) Downhole pressure

(b) Choke flow

(c) Pump flows

(d) Choke opening

Figure 3.3: Connection response of original MPD system in calm conditions

3.1.4 Connection Response in Rough Conditions

A connection done in rough conditions changes the response quite a lot and can be seen in Fig. 3.4. Just like in calm conditions (Fig. 3.3), we can see that the mean choke opening, without CCS implemented, decreases due to the decreased flow during the connection. And again, we see that with CCS, the opening is still operating around 50%. However, we now see major oscillations in both downhole pressure, choke flow and choke opening. The downhole pressure and choke flow oscillations (Fig. 3.4a and 3.4b) are results of the heaving motion of the rig and drill string. The oscillations seen in the choke opening (Fig. 3.4d) is the result of the MPD system trying to compensate for the oscillations in downhole pressure and/or choke flow. However, as will become more obvious in Section 3.2, manipulation of the choke opening does not seem to compensate for the downhole pressure oscillations at all. The MPD system is still able to keep the downhole pressure around its desired value though, but the oscillations in the choke opening are unnecessary and should be prevented to decrease wear and tear.

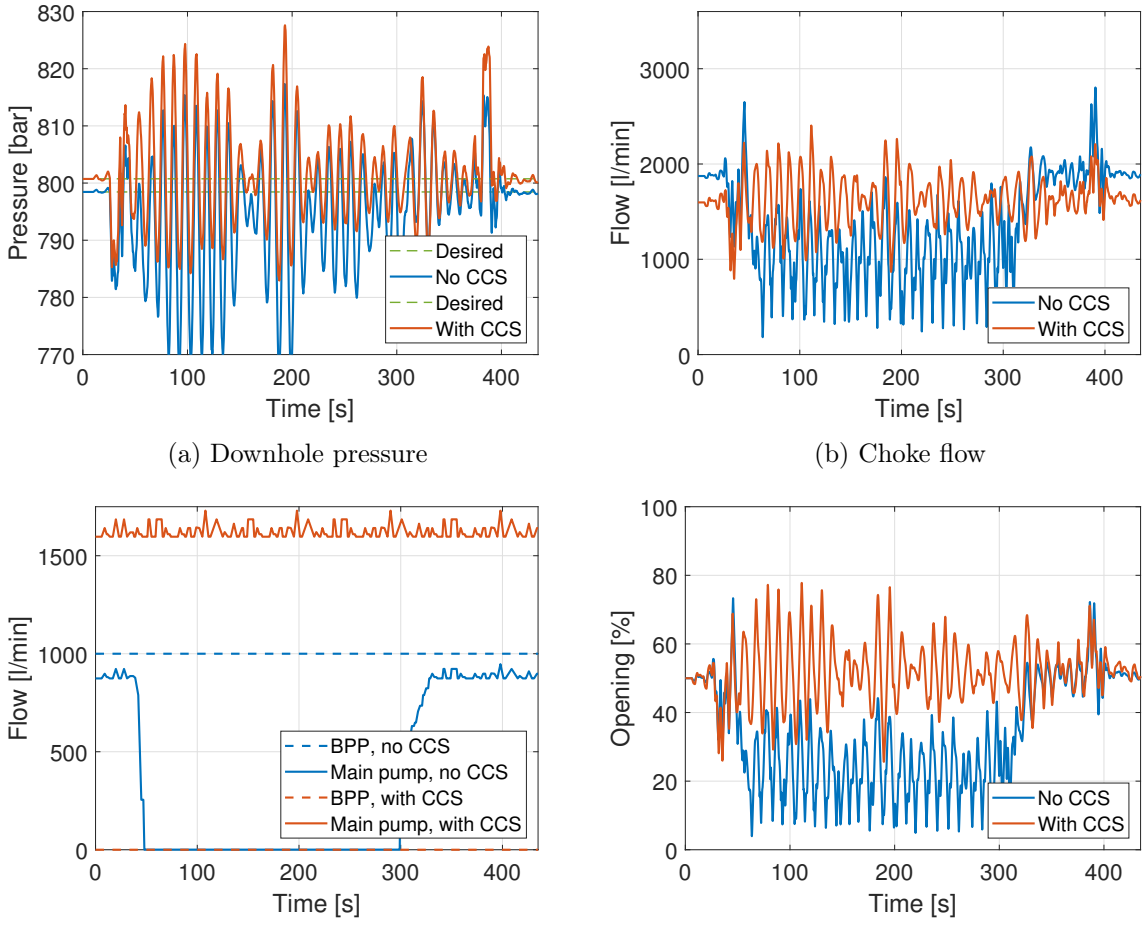

(a) Downhole pressure

(b) Choke flow

(c) Pump flows

(d) Choke opening

Figure 3.4: Connection response of original MPD system in rough conditions

3.2 Amplification of Heave-induced Oscillations

Simulations have been done to try to show the effect of the delay described in Section 2.3 in a realistic manner. The example mentioned there was simplified quite a lot just to illustrate the effect, and the effect is thus not expected to be quite as obvious in the following simulations. In real wells, there are other conditions that need to be accounted for. For example, the friction in a well, especially in a non-vertical well, has a huge effect on the heave induced pressure changes downhole as the well gets deeper. In a shallow well, or a perfectly vertical one, where friction is close to non-existent, the drill bit moves close to as much as the drilling rig does. This means that the heave motion topside is almost identical to the one downhole. As the well gets deeper though, taking turns and maybe even becoming horizontal, stiction plays a huge role. Even though the rig is moving vertically, the motion might not be enough to overcome the stiction forces downhole, and the drill bit may not move at all. This makes it harder to produce good simulations that show the effect of delay consistently.

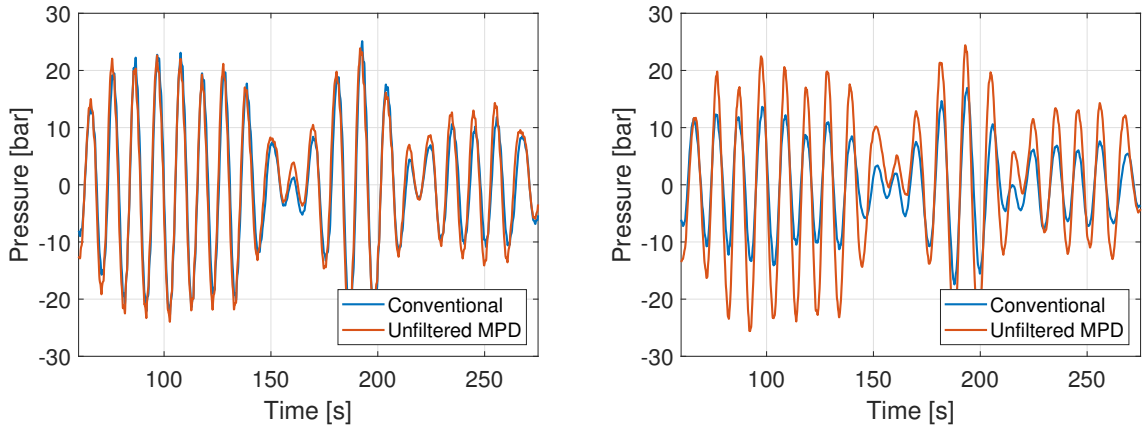

(a) Connection at 3000m

(b) Connection at 4000m

Figure 3.5: Comparison of downhole pressure variations at different measured depths (MDs)

In Fig. 3.5 there is a comparison between conventional drilling and unfiltered MPD where a snippet of the downhole pressure oscillations from simulations of a connection done at two different depths are plotted. At 3000 meters MD we see that conventional drilling has the largest amplitude and unfiltered MPD decreases the pressure variations. At 4000 meters the opposite is evident. We can clearly see that unfiltered MPD does have an effect on the significance of pressure variations, in some cases for the better, but in some for worse. To try to find a way of showing how conventional drilling compares more clearly, several more simulations were done at depths ranging from 1000 to 4000 meters.

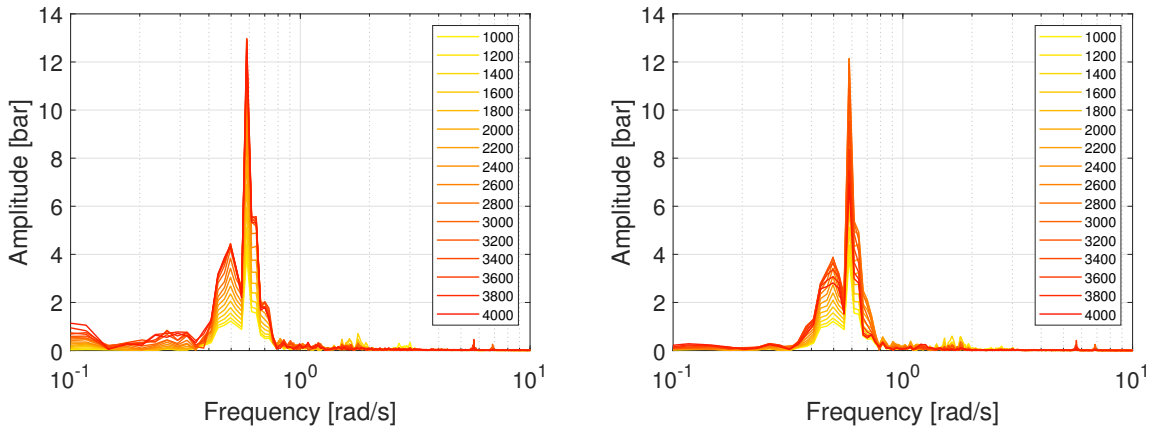

(a) Conventional Drilling (b) Unfiltered MPD

Figure 3.6: Discrete Fourier Transform of downhole pressures

An easy way to compare the signals is obviously to look at the amplitude of both conventional drilling and unfiltered MPD. The Fourier Transform (FT) was used to extract information about the frequencies and amplitudes in the signals to easier identify the differences, rather than comparing a lot of plots with each other visually. Fig. 3.6 shows the Discrete Fourier Transform (DFT) of the downhole pressure of every simulation, both with conventional drilling and the original unfiltered MPD system. We can see that the most evident frequency is around 0.58 rad/s, i.e. waves with a period of about 10.8 seconds. It still is not obvious, though, if MPD actually is amplifying the pressure variations, because a quick glance at the figure shows greater pressure variations with conventional drilling.

If we plot the most evident amplitude as a function of depth we get a better look at the differences, see Fig. 3.7. Here we can see that the performance of both MPD and conventional drilling is almost identical from 1000 to about 2500 meters TVD. However, the amplitude seems to decrease after 2800 meters when using conventional drilling, and continue to increase with MPD implemented.

I had a hard time trying to figure out why conventional drilling has decreasing amplitudes as the depth exceeds 2800 meters TVD, but no matter why or how, though, the performance of the MPD system is what is important. Although MPD did not seem to have much effect on the downhole pressure oscillations at shallow depths, simulations showed a choke opening trying to compensate for said oscillations without much luck. And even though the system tried to compensate for the oscillations, we see that it starts amplifying them, rather than compensate for them, at depths exceeding 3000 meters TVD. Amplification of the downhole pressure is, of course, undesirable, but so is the fact that the choke opening is operated unnecessarily. If there

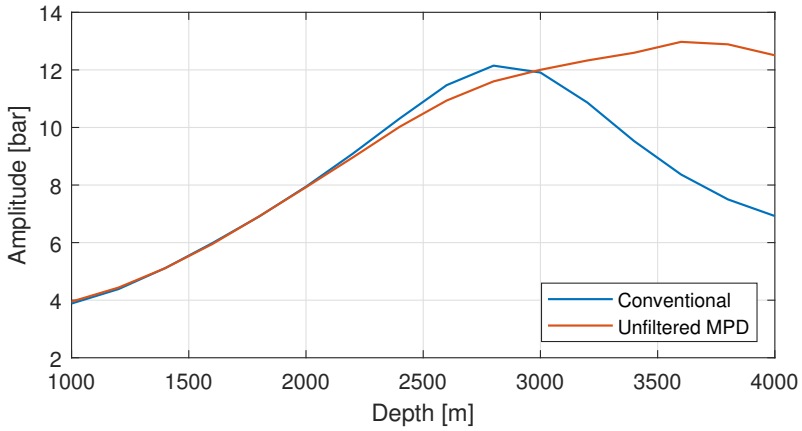

Figure 3.7: Comparison of downhole pressure amplitudes at different measured depths (MDs)

is no positive effect of controlling the valve, it should be prevented to minimize any wear and tear and reduce maintenance costs.

This means that we do not want the MPD system to intervene at all when exposed to heave-induced oscillations. Preferably, we would like to see a topside choke opening that more or less stays stationary during a connection, and a response during connections that is similar to the conventional one, where there is no topside choke, or it has a constant opening. In other words, the downhole pressure amplitudes for MPD in Fig. 3.7 should follow the conventional amplitudes. To manage this, some kind of filter must be implemented in the MPD control system to attenuate the frequencies we do not want to control.

4 Filter Design

As mentioned previously, a filter should be implemented in the MPD control system, not only to try eliminating the amplification effect that the original MPD system has on the downhole pressure oscillations at some depths but also to reduce any wear and tear on the choke valve, as this is unnecessarily operated. To find what kind of filter we need to implement, we should first figure out which frequencies we need and which we should get rid of. The frequencies we want to keep are obviously the low frequencies, as the MPD system is designed for slowly varying conditions, but where should the boundary that defines which low frequencies we want to keep be? To find this boundary we need to analyze the simulation results.

4.1 Frequency Analysis

The obvious frequencies to remove are the ones that exist in the heaving motion of the drilling rig. Heave is what creates the pressure variations downhole during a connection, meaning that the frequencies found in the downhole pressure during such a connection should correspond to those found in the heaving motion, and we should, therefore, start with analyzing this.

The heaving motion is a result of the drilling rig being influenced by weather disturbances, mainly ocean waves, tides, and currents. Ocean waves are probably the most influential disturbance for short-term connections though. A simple DFT analysis of the heave data used in the simulations show that the peak frequency for the rig heave is about 0.58 rad/s, which corresponds to a mean wave period of about 10.8 seconds and concurs with the frequencies found in the downhole pressure during a connection in Section 3.2. Fig. 4.1 shows that the frequency spectrum of the heave data ranges from about 0.4 to 0.8 rad/s, which corresponds to waves with periods ranging from about 8 to 16 seconds. This is just the DFT results of a single sample of heave data, though, and more data should be gathered to do a proper frequency analysis. Gathering data in different weather conditions will probably reveal a wider frequency range than what is found in Fig. 4.1b, and may lead to consideration of cascaded filters in a real implementation. The heave data used here, however, actually fits very well with the Torsethaugen spectrum, a wave spectrum developed with data from the North Sea [21].

The fact that the same frequencies are found in both the heave and downhole pressure data confirms that heave is the reason behind the pressure oscillations during a connection and that these are the frequencies we need to focus on filtering.

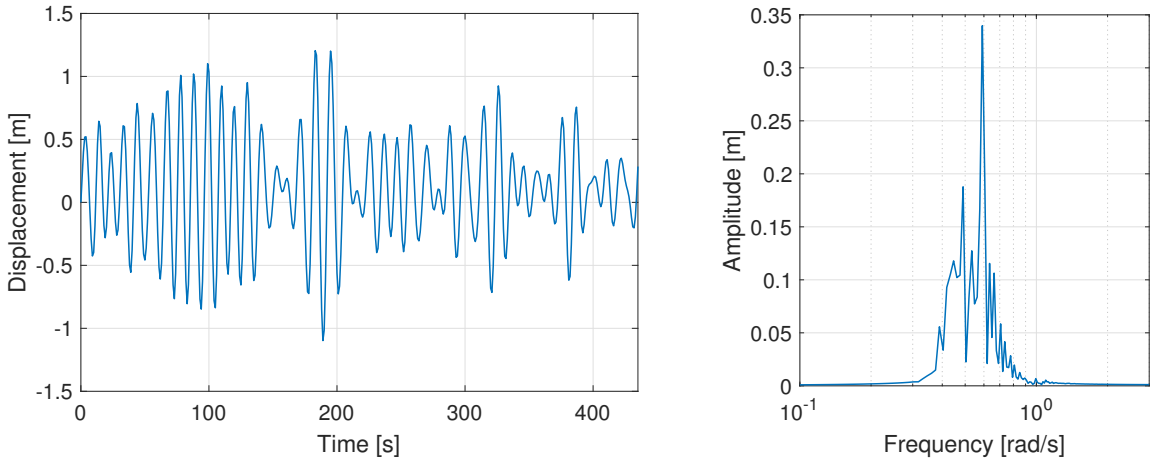

(a) Heaving motion (b) Discrete fourier transform

Figure 4.1: Heave data analysis

4.2 Implementation

Different filtering strategies to experiment with were presented in Section 2.4, each with their own abilities and applications, but first, we must identify which states that should be filtered.

The desired choke-flow in (3.2) is the main control law in the MPD system, where (3.1) determines the desired choke pressure based on a desired downhole pressure. One could imagine that if we were to filter the desired choke flow, to keep this constant, the choke opening should also stay constant. However, we need to remember that the downhole pressure oscillations transfer topside, meaning that the topside choke pressure is also affected by the heaving motion of the drill string. Even if the desired choke flow was constant, the choke opening would still oscillate as it is a function of the choke pressure (3.3). It is exactly the choke opening that we wish was unaffected by the pressure oscillations, but, the choke opening is nonlinear with respect to the choke pressure and filtering the desired opening with a linear filter can then give uncertain results, especially if the differential choke pressure creeps toward zero. A better approach should be to filter both the desired choke flow and the differential pressure separately, as shown in the block diagram in Fig. 4.2, making both more or less constant and resulting in a more stationary choke opening.

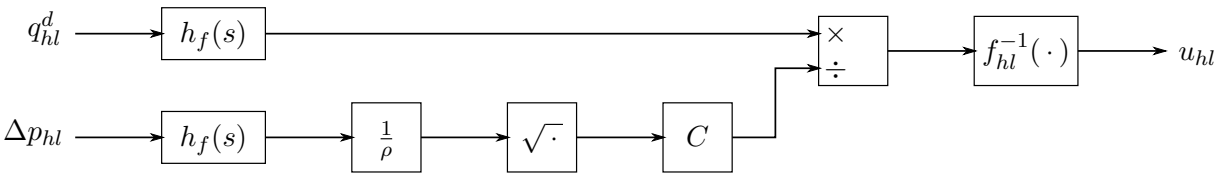

Figure 4.2: MPD filter implementation

Several different filters were implemented and tested in the simulator, including both simple

low-pass filters and notch filters of different orders. The one that showed the best results was based on the filter suggested by Michael J. Grimble and Michael A. Johnson defined in (2.14) [20]. The time constant of the low-pass filter was, however, chosen as $\tau = 0.8$ s instead of 0.1 s. The same kind of filter was used on both the desired choke flow and the choke's differential pressure.

4.2.1 Simulations

The simulation results from both a setpoint step response and a pump flow ramp response of the MPD system with the filter defined in (2.14) implemented is plotted in Fig. 4.3. Both responses is with CCS implemented, see Appendix C to see the responses without CCS. In comparison with the unfiltered responses in Fig. 3.1 and 3.2 we see that the responses are fairly similar, but with less noise. The noisy pump pressure in Fig. 4.3c no longer seems to influence the choke opening in Fig. 4.3d. We do, however, see a difference in the setpoint step response, as the downhole pressure now has a couple of overshoots before it stabilizes. This is probably due to the phase lag that the filter introduces. The overshoot can be reduced at the expense of a slower response and slightly worsened pump flow ramp response by decreasing the k_p gain in (3.2), but should not necessarily be an issue if we assume that big steps in desired downhole pressure are rare.

The effects of the filter are much more obvious if we take a look at the simulation results of a connection done in rough weather. We saw in Section 3.1.4 and Fig. 3.4 that the original MPD system unsuccessfully tried to compensate for the heave induced downhole pressure oscillations during a connection. In Fig. 4.4 we see a comparison of a connection done in rough conditions with and without the filter implemented. With the filter implemented the MPD system no longer tries to compensate for the induced pressure oscillations, as seen from the topside choke opening in Fig. 4.4d that now stays close to constant throughout the connection. As a result, we actually see that the downhole pressure oscillations are reduced slightly.

Several more simulations were run to analyze the performance of the filter as well as compare it with conventional drilling and the unfiltered performance of the MPD system. The simulations were again of connections at depths ranging from 1000 to 4000 meters, just like the ones ran in Section 3.2. By again plotting the most evident amplitude as a function of depth we can see how the different systems compare while performing connections, see Fig. 4.5. We see that the implemented filter helps reduce the downhole pressure oscillations at all depths compared to the

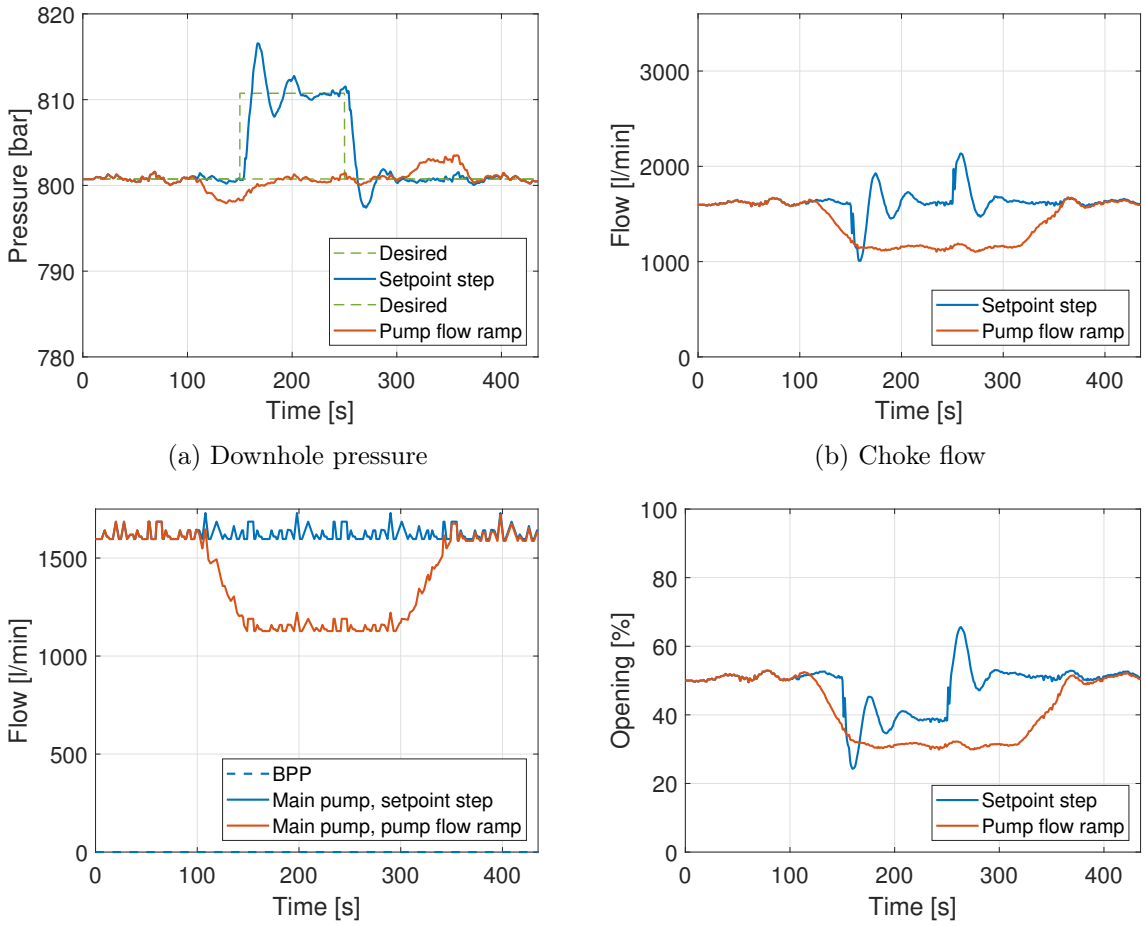

(a) Downhole pressure

(b) Choke flow

(c) Pump flows

(d) Choke opening

Figure 4.3: Step and ramp responses of MPD system with filter

unfiltered MPD performance. CCSs help reduce the oscillations even further. However, in all cases, the MPD system still seems to amplify the oscillations somewhat at great depths when compared to conventional drillings. This is not as we thought in Section 3.2, where we expected the filtered MPD system to follow the performance of conventional drilling. Although our expectations were wrong, the filter still has positive effects on the downhole pressure variations when compared to the unfiltered system, while almost eliminating any manipulation of the topside choke opening. This should be a good enough reason to implement such a filter if MPD were to be considered used on a floating drilling rig.

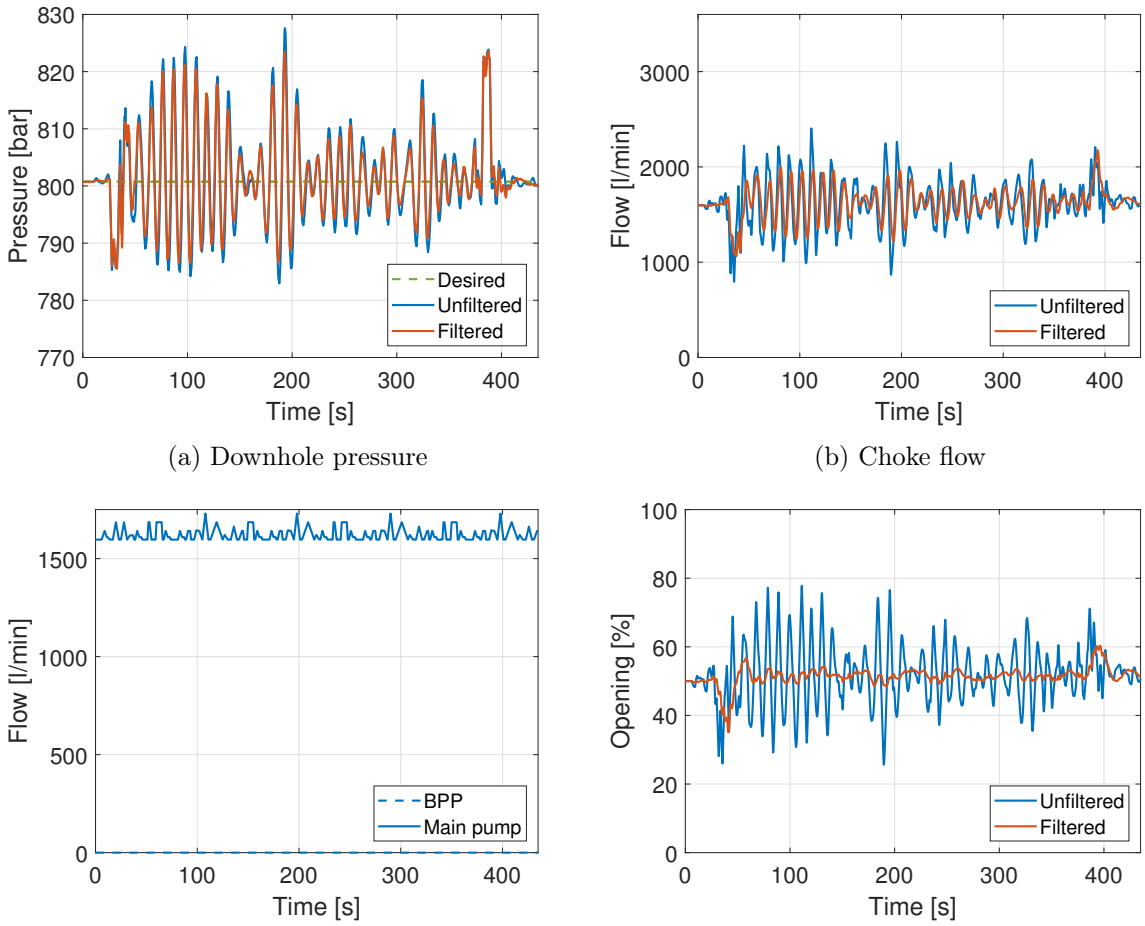

(a) Downhole pressure

(b) Choke flow

(c) Pump flows

(d) Choke opening

Figure 4.4: Connection response of MPD system with and without filter in rough conditions

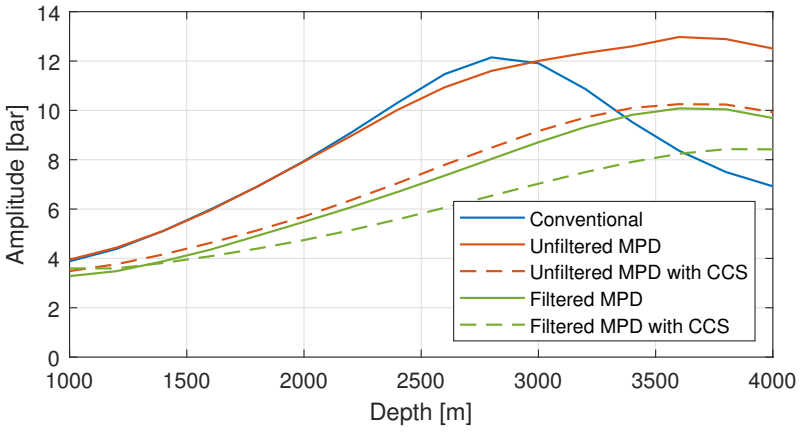

Figure 4.5: Comparison of downhole pressure amplitudes with and without filter implemented

5 HeaveLock™ Controller

We have seen that the MPD system presented in Section 3 together with the filter designed in Section 4 performs well with changes in both setpoint and pump flow. We have also seen that it no longer tries to compensate for heave-induced pressure variations during connections, with positive effects on downhole pressure oscillations and less manipulation of the choke opening, reducing wear and tear and possibly future maintenance costs. We will from now on assume that this system performs well enough for it to be used on floating drilling rigs, so we can begin studying if and how HeaveLock™ will affect its performance or vice versa. Let us first take a look at HeaveLock™'s controller.

The original code for the HeaveLock™ controller in HeaveSim allows for testing of both feedback control, feed-forward control or both. Measurements like inlet, outlet and annulus pressure outside of HeaveLock™, as well as mud-flow and bit velocity, are available for testing of controller strategies. The feedback controller is a tunable PID-controller that uses the annulus pressure outside of HeaveLock™ as feedback with the option of a tracking value that is the difference between actual and desired flow through HeaveLock™. The feed-forward controller is implemented as a state-space system, where the inputs to the system are the measurements mentioned above, and the number of states and outputs are defined by the system matrices dimensions, i.e.

$$\dot{\boldsymbol{x}} = \boldsymbol{A}\boldsymbol{x} + \boldsymbol{B}\boldsymbol{u} \tag{5.1a}$$

$$\boldsymbol{y} = \boldsymbol{C}\boldsymbol{x} + \boldsymbol{D}\boldsymbol{u}, \tag{5.1b}$$

where $\boldsymbol{x} = [x_1 \ \dots \ x_n]^\top$, $\boldsymbol{u} = [u_1 \ \dots \ u_r]^\top = [p_{hl0} \ p_{hl} \ p_{hl,a} \ q_{hl} \ v_b \ e_{p_{hl,a}}]^\top$ and $\boldsymbol{y} = [y_1 \ y_2]^\top$. The first output, y_1, is reserved as a potentially filtered bit velocity, while y_2 is an optional feed-forward term. The desired flow through HeaveLock™ is calculated as

$$q_{hl}^d = q_{hl}^{nom} - A_d y_1 + y_2 + y_{pid}, \tag{5.2}$$

where y_1 and y_2 are outputs from the state-space system, y_{pid} is the output from the feedback PID-controller, q_{hl}^{nom} is the nominal flow through HeaveLock™ and A_d is the area of the volume that the drill string displaces downhole. The block diagram in Fig. 5.1 also shows the needed inputs and the output of the controller. The nominal flow is the measured flow through HeaveLock™ at activation and acts as an estimate of the pump flow, i.e. $q_{hl}^{nom} = \hat{q}_p$. Once the desired

flow has been calculated, the control input u_{hl} for HeaveLock™ is found with

$$u_{hl} = f_{hl}^{-1}\left(\frac{q_{hl}^{d}}{C\sqrt{\frac{\Delta p_{hl}}{\rho}}}\right), \tag{5.3}$$

where $f_{hl}(\cdot)$ is HeaveLock™'s characteristic function that ranges from 0 to 1, C is its valve constant, Δp_{hl} is the differential pressure over the valve and ρ is the mud density. The function argument in (5.3) is built into the f_{hl}^{-1} function block in the block diagram in Fig. 5.1, see Fig. 4.2 for how.

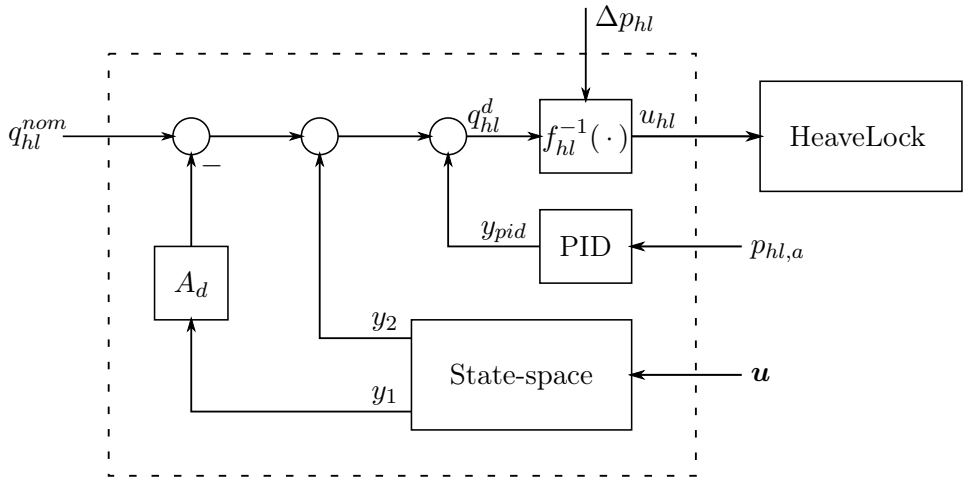

Figure 5.1: HeaveLock™ controller

5.1 Original Strategy

The original control strategy for HeaveLock™ is a simple one. The volume change due to the drill string's motion is to be dealt with by controlling the mud-flow through the drill string. I.e., if the drill string is heaving upwards, hypothetically creating an empty volume beneath it, this volume is to be filled with an increased mud flow through the drill string, see Fig. 5.2a. The opposite applies when the drill string is heaving downwards. The displaced volume is to be accounted for by restricting the mud-flow, see Fig. 5.2b. This control strategy could, in theory, result in a constant annular flow, and should decrease surge and swab pressures. The volumetric flow rate due to drill string motion (or the desired change in flow rate) is simply computed as

$$\dot{V} = \Delta q_{hl} = A_d v_b, \tag{5.4}$$

where A_d is the area of the displaced volume and v_b is the axial velocity of the bit relative to the well. The desired flow through HeaveLock™ is then

$$q_{hl}^d = \hat{q}_p - \Delta q_{hl} = \hat{q}_p - A_d v_b, \tag{5.5}$$

where $\hat{q}_p$ is an estimate of the main pump flow and positive bit velocity is defined in the downwards axial direction.

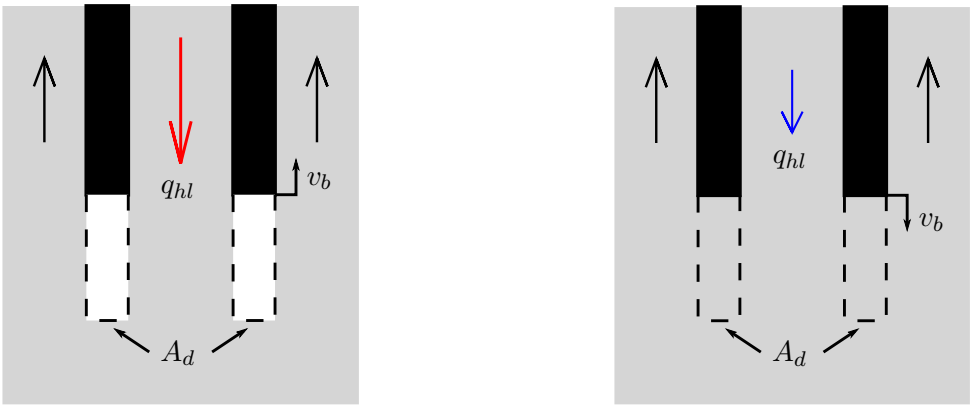

(a) Heaving upwards creates an imaginary void (b) Heaving downwards will displace a volume

Figure 5.2: Volume change due to axial bit motion

5.1.1 Implementation and Performance

Unfiltered Velocity

To implement the unfiltered case of the simple controller given in (5.5) in HeaveSim, the state-space system matrices can be written as follows

$$\boldsymbol{A} = a = 0 \qquad \boldsymbol{B} = \boldsymbol{b}^\top = [0\ 0\ 0\ 0\ 0\ 0] \tag{5.6a}$$

$$\boldsymbol{C} = c = 0 \qquad \boldsymbol{D} = \boldsymbol{d}^\top = [0\ 0\ 0\ 0\ 1\ 0], \tag{5.6b}$$

and the PID-controller gains must all be set to zero. This simply translates to $q_{hl}^d = \hat{q}_p - A_d y_1$, where $y_1 = v_b$ is just the measured bit velocity. Fig. 5.3 shows the results from a simulation of a connection where HeaveLock™'s original control strategy and unfiltered velocity was implemented. The figure shows a comparison of using the full bit velocity and scaling it down by half. For some reason, HeaveLock™'s controller seem to pick up on a higher frequency which destabilizes the system when the full bit velocity is used. By reducing the velocity, the system is stable and HeaveLock™ is able to reduce the downhole pressure oscillations.

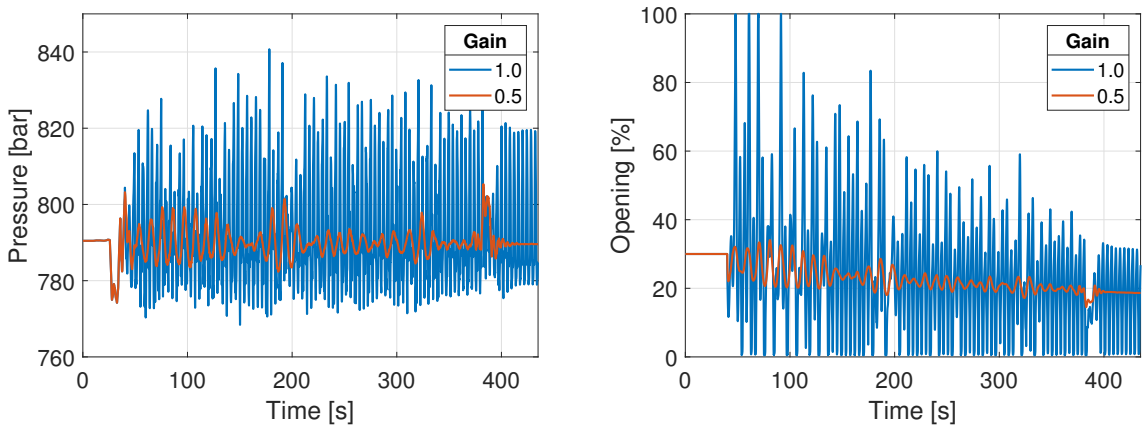

(a) Downhole pressure (b) HeaveLock™'s opening

Figure 5.3: Original HeaveLock™ controller with unfiltered velocity

Filtered Velocity

If we would like to use a filtered bit velocity instead of an unfiltered one, this is also possible. Like the results in Fig. 5.3 suggests, the filter should attenuate higher frequencies than those represented in the heaving motion. Removing low frequencies may also be a good idea to prevent HeaveLock™ trying to compensate for motion related to hoisting or lowering of the drill string. Hoisting the drill string with a constant velocity, for example, will make the desired flow in (5.5) constantly higher than the pump flow. This will then increase HeaveLock™'s opening, as seen from (5.3), which will cause the differential pressure to decrease, which in turn will lead to an even greater opening. This means that, without attenuating low frequencies, long-lasting motion in one direction may drive HeaveLock™ to saturation.

A band-pass filter with a passband containing the most evident frequencies in the heave motion will attenuate the lower and higher frequencies while leaving the frequencies of heave untouched. HeaveLock™ should then focus on compensating for the heave-induced oscillations and nothing more. A state-space system of the fourth order Butterworth band-pass filter defined in (2.10) can be found by using the observable canonical form defined in (2.22), i.e.

$$\dot{\boldsymbol{x}} = \begin{bmatrix} -\sqrt{2}\omega_{bw} & 1 & 0 & 0 \\ -(2\omega_c^2 + \omega_{bw}^2) & 0 & 1 & 0 \\ -\sqrt{2}\omega_{bw}\omega_c^2 & 0 & 0 & 1 \\ -w_c^4 & 0 & 0 & 0 \end{bmatrix} \boldsymbol{x} + \begin{bmatrix} 0 \\ \omega_{bw}^2 \\ 0 \\ 0 \end{bmatrix} u \tag{5.7a}$$

$$y = \begin{bmatrix} 1 & 0 & 0 & 0 \end{bmatrix} \boldsymbol{x}. \tag{5.7b}$$

By replacing u with the bit velocity that we want to be filtering, v_b, HeaveLock™'s state-space matrices can be written as

$$\boldsymbol{A} = \begin{bmatrix} -\sqrt{2}\omega_{bw} & 1 & 0 & 0 \\ -(2\omega_c^2 + \omega_{bw}^2) & 0 & 1 & 0 \\ -\sqrt{2}\omega_{bw}\omega_c^2 & 0 & 0 & 1 \\ -w_c^4 & 0 & 0 & 0 \end{bmatrix} \qquad \boldsymbol{B} = \begin{bmatrix} 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & \omega_{bw}^2 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 \end{bmatrix} \tag{5.8a}$$

$$\boldsymbol{C} = \boldsymbol{c}^\top = [1\ 0\ 0\ 0] \qquad \boldsymbol{D} = d = 0, \tag{5.8b}$$

which translates to $q_{hl}^d = \hat{q}_p - A_d v_{b,f}$, where $v_{b,f}$ is the filtered version of the bit velocity v_b. Fig. 5.4 shows the results from a simulation of a connection where the filter defined above was implemented in HeaveLock™'s original control strategy. The bandwidth and center frequency of the filter was chosen as 0.6 and 0.63 rad/s, respectively. Again, the figure shows a comparison of using the full bit velocity and a scaled down version. The system is now stable in both cases, and we see that the full velocity gives the best results, i.e. the lowest variation in downhole pressure.

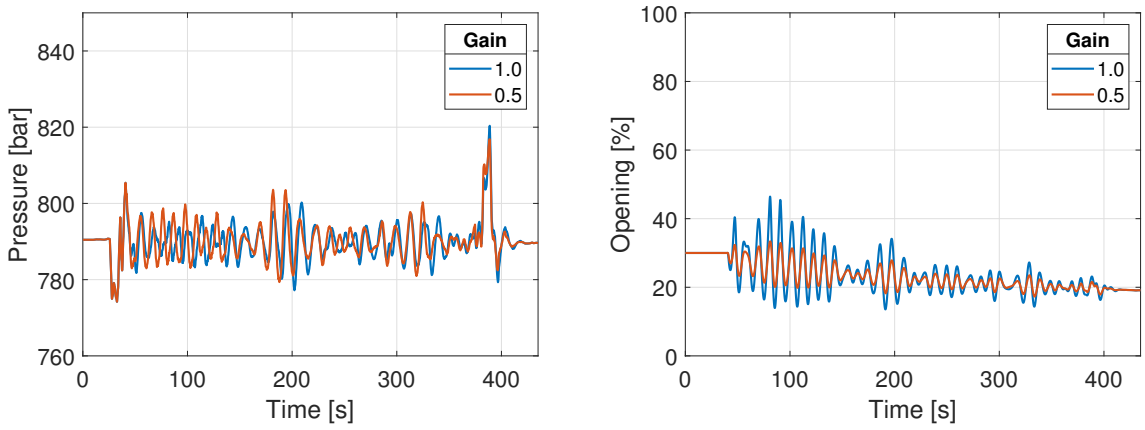

(a) Downhole pressure

(b) HeaveLock™'s opening

Figure 5.4: Original HeaveLock™ controller with band-pass filtered velocity

5.2 Drift Problem

The original control strategy has a problem with drift. As we can see from Fig. 5.5, when time goes on, the opening of HeaveLock™ slowly decreases, increasing the differential pressure over it, which also means an increasing pump pressure. The nature of this happening is the fact that we only try to control the transient pressure downstream from HeaveLock™, i.e. the downhole

pressure variations, and that we have not given HeaveLock™'s opening or upstream conditions much consideration.

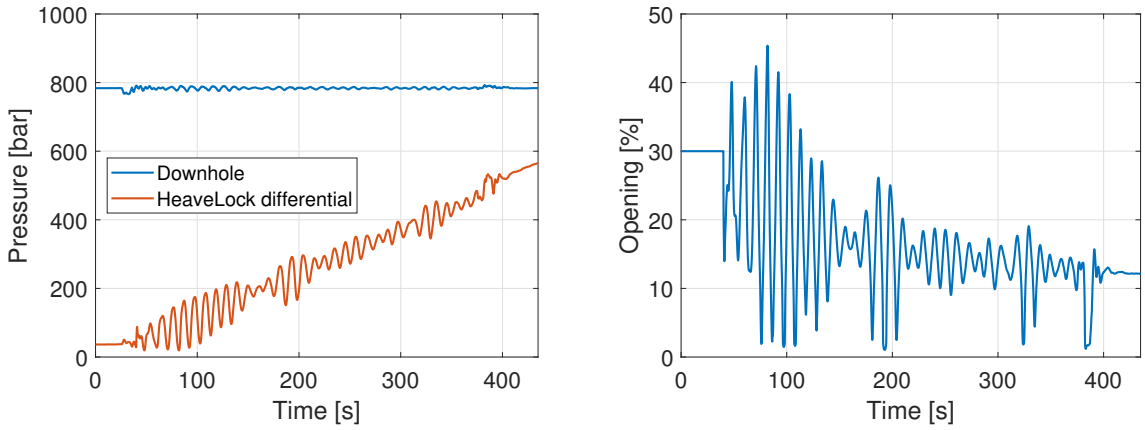

(a) Pressures (b) HeaveLock™'s opening

Figure 5.5: Drift when HeaveLock™ is initiated with low pump flow

The problem with drift, in this case, is that HeaveLock™'s differential pressure and the pump pressure may exceed their recommended limits, which could lead to other problems. The differential pressure over a valve is also much more sensitive to changes when the valve's opening is small. Also, if HeaveLock™'s opening were to increase, the built-up pressure in the drill string would decrease, weakening HeaveLock™'s performance. Drift may also drive HeaveLock™ closer to saturation, in effect losing its controllability.

The reason for why drift occurs has to do with how HeaveLock™ is initiated. Its controller computes an estimate of the pump flow at initiation based on inlet and outlet pressures and assumes that this value is the correct pump flow to be used in (5.2) throughout HeaveLock™'s operating window. If the pump flow is initiated too low, as is the case in Fig. 5.5, the desired flow will be less than optimal, resulting in a smaller HeaveLock™ opening and a greater pump pressure. In Fig. 5.6 we see how HeaveLock™ will drift if the pump flow is initiated too high. The desired flow will then always be higher than optimal, resulting in a larger opening. In this case, it even drives HeaveLock™ to saturation, and we can see that the differential pressure is slowly eliminated, meaning that HeaveLock™'s performance will be weakened due to the lack of built-up pressure in the drill string.

Since we aim to study if HeaveLock™ will affect the MPD system's performance and vice versa, it is important that the controller we implement is a realistic one. It is not likely that the controller that will be commercialized will have any drift, and we should, therefore, find a solution to eliminate said drift before we study how HeaveLock™ and the MPD system will interact.

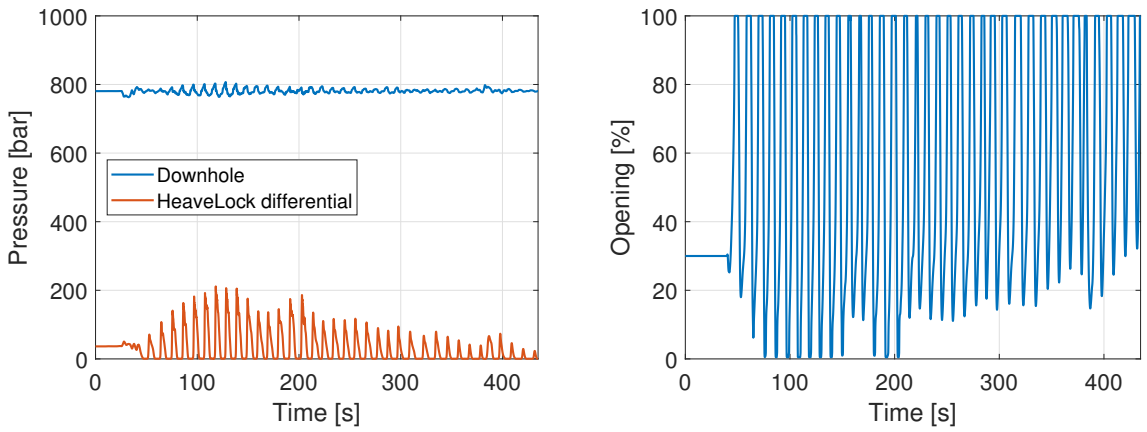

(a) Pressures (b) HeaveLock™'s opening

Figure 5.6: Drift when HeaveLock™ is initiated with high pump flow

5.3 Eliminating Drift

The pressure at a point in a pipeline is determined by downstream conditions like a pressure reference, the pressure drops due to friction in a restriction or along a surface, and the hydrostatic pressure due to elevation. In this case, where the flow is one-directional from the main pump to the topside choke, the reference pressure will be the atmospheric pressure at the choke's outlet and the stationary downhole pressure is only determined by the pressure drop over the choke (back-pressure), the hydrostatic pressure at the given TVD and the annular friction loss, see (2.1). This means that stationary changes in HeaveLock™'s opening, and its respective differential pressure, will only have stationary effects on the upstream pressures from HeaveLock™, i.e. the pressure inside the drill string and the pump pressure and not the downhole pressure. Upstream valves like HeaveLock™ can, of course, have transient effects on the downstream pressure, meaning that manipulation of HeaveLock™'s opening will still have transient effects on the downhole pressure.

It then follows that the stationary downhole pressure is independent of the stationary differential pressure over HeaveLock™. We are simply allowing HeaveLock™ to drift by not caring about its opening or upstream conditions. If we keep an eye on these, we can correct for any occurring drift and eliminate the problem.

The easiest way to eliminate drift is to punish deviation from some desired value. In this case, it is the drift in HeaveLock™'s opening that causes the problem. If we had a nominal opening for HeaveLock™, we could in theory just punish any deviation from this nominal value with an integrator term in HeaveLock™'s controller, and the problem should be solved. But the

question is: how do we find such a nominal value for its opening? We do not want a nominal opening that is less than necessary, since, as mentioned earlier, the differential pressure over a valve is more sensitive if the valve is operating around a small opening. The valve will also be exposed to greater forces with a small opening, requiring more frequent maintenance. Still, the nominal opening must be small enough to build up the necessary pressure in the drill string for HeaveLock™ to function as designed.

The desired opening of HeaveLock™ can be found with equation (5.3), which is a function of the desired flow and the differential pressure over the valve. To find a nominal opening, we can use the same equation, but we need a nominal value for both the flow and differential pressure

$$u_{hl}^{nom} = f_{hl}^{-1}\left(\frac{q_{hl}^{nom}}{C\sqrt{\frac{\Delta p^{nom}}{\rho_d}}}\right). \tag{5.9}$$

The flow is easy, as in most cases when HeaveLock™ is active, the main pump will deliver a constant flow. The nominal flow should therefore be equal to the main pump's flow, i.e. $q_{hl}^{nom} = q_p$. The exact pump flow cannot be measured downhole, though, and the flow through HeaveLock™ at activation must act as an estimate of the pump flow. The nominal differential pressure, however, must be found by other means. The differential pressure over HeaveLock™ determines how much pressure that is built up in the drill string above it. Just putting some very high value as the nominal value would ensure that there is enough pressure to work with at all times, but this would again result in an unnecessary small opening. We need to find a nominal differential pressure that ensures good performance without risking excessive maintenance.

5.3.1 Exploiting Knowledge of Bulk Modulus

By knowing some data of the drill string and mud, it is possible to calculate the compression of a volume due to pressure. A fluid's bulk modulus, β, is defined as [22]

$$\frac{d\rho}{\rho} = \frac{dp}{\beta} \qquad \Longleftrightarrow \qquad \frac{dV}{V} = -\frac{dp}{\beta}, \tag{5.10}$$

where ρ is the mud density and p is the pressure inside the fluid volume V. From this we have that a positive change in pressure results in a negative change in volume, see Fig. 5.7.

We would like to have built up enough pressure, such that the compressed volume in the drill string is enough to fill whatever the needed volume beneath the bit is when it expands. If the

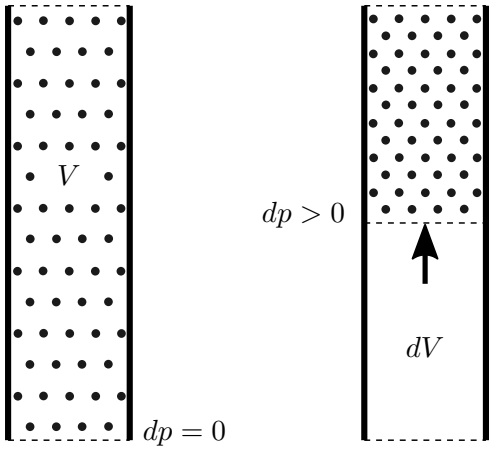

Figure 5.7: Bulk modulus

drill string is pulled up by a length l, the volume to fill beneath the bit would be the length multiplied with some area A. The change in volume due to compression inside the drill string should be equal to this value, i.e. the desired change in volume is

$$dV^d = A_d l^d, \tag{5.11}$$

where A_d is the displaced area (see Fig. 5.2) and l^d is the desired length of axial motion that HeaveLock™ should be able to control, e.g. the peak-to-peak value of the heaving motion. The desired differential pressure over HeaveLock™ to compress the volume by such an amount is therefore

$$dp^d = \frac{\beta_d}{V_d} dV^d = \frac{\beta_d}{V_d} A_d l^d, \tag{5.12}$$

where β_d is the bulk modulus of the fluid and V_d is the volume inside the drill string. The problem is that the volume inside the drill string varies with every added stand, and is forever increasing during the drilling of a well. At a given moment of time, it is difficult for HeaveLock™ to know the exact volume. However, the hydraulic area A_h inside the drill string is known and does not change by much, so the problem reduces to not knowing the length of the drill string l_{ds}. The desired differential pressure is therefore dependent on the length of the drill string, i.e.

$$dp^d = \frac{A_d \beta_d}{A_h l_{ds}} l^d, \tag{5.13}$$

but without somehow knowing the length of the drill string, l_{ds}, this approach achieves nothing.

5.3.2 Estimating the Length of the Drill String

Minimum Length or True Vertical Depth

The downhole pressure p_{dh} at any given TVD is given by (2.1). The hydrostatic pressure is then given as

$$p_{h,a} = p_{dh} - p_c - p_{f,a}, \tag{5.14}$$

where the choke pressure and friction loss are both positive and may vary. We can use the hydrostatic pressure $p_{h,a} = \rho_a g h$ as a basis for finding the minimum length of the drill string, i.e.

$$l_{ds}^{min} = \frac{p_{h,a}}{\rho_a g}, \tag{5.15}$$

where ρ_a is the mud density and g is the gravitational acceleration. The problem is that if we were to measure the downhole pressure to calculate the hydrostatic pressure, the measurement would be quite uncertain because of varying friction loss due to varying flow rate, and the unknown amount of back-pressure applied. The mud density ρ_a can also change during drilling, especially if mud engineers are manipulating it. However, if we know the maximum values for applied back-pressure (choke pressure) and friction loss, p_c^{max} and p_f^{max}, we can utilize this and write the hydrostatic pressure in (5.14) as

$$p_{h,a} = p_{dh} - \frac{1}{2}(p_c^{max} + p_f^{max}) \pm \frac{1}{2}(p_c^{max} + p_f^{max}) = \hat{p}_{h,a} \pm \Delta p_{h,a}, \tag{5.16}$$

where $\hat{p}_{h,a} = p_{dh} - \frac{1}{2}(p_c^{max} + p_f^{max})$ is the best estimate of the hydrostatic pressure, $\Delta p_{h,a} = \frac{1}{2}(p_c^{max} + p_f^{max})$ is the uncertainty and p_{dh} is the measured downhole pressure. We can also write the mud density and gravitational acceleration with their uncertainties

$$\rho_a = \hat{\rho}_a \pm \Delta\rho_a \tag{5.17}$$

$$g = \hat{g} \pm \Delta g, \tag{5.18}$$

and with this information we can calculate the relative minimum depth uncertainty, based on downhole measurements

$$\frac{\Delta l_{ds}^{min}}{\hat{l}_{ds}^{min}} = \frac{\Delta p_{h,a}}{\hat{p}_{h,a}} + \frac{\Delta\rho_a}{\hat{\rho}_a} + \frac{\Delta g}{\hat{g}}, \tag{5.19}$$

where

$$\hat{l}_{ds}^{min} = \frac{\hat{p}_{h,a}}{\hat{\rho}_a \hat{g}} = \frac{p_{dh} - \frac{1}{2}(p_b^{max} + p_f^{max})}{\hat{\rho}_a \hat{g}}, \tag{5.20}$$

and p_{dh} is the downhole measurement at HeaveLock™ ($p_{hl,a}$). Assuming we also have information about the displaced area A_d, the drill string's hydraulic area A_h and annular mud bulk modulus β_a, we can find the best estimate of the maximum needed differential pressure over HeaveLock™ with (5.13)

$$\hat{dp}^{max} = \frac{\hat{A}_d \hat{\beta}_d}{\hat{A}_h \hat{l}_{ds}^{min}} l^d, \tag{5.21}$$

with its relative uncertainty

$$\frac{\Delta dp^{max}}{\hat{dp}^{max}} = \frac{\Delta A_d}{\hat{A}_d} + \frac{\Delta \beta_d}{\hat{\beta}_d} + \frac{\Delta A_h}{\hat{A}_h} + \frac{\Delta l_{ds}^{min}}{\hat{l}_{ds}^{min}}. \tag{5.22}$$

To further ensure that the maximum differential pressure is sufficient, it could be wise to add the absolute uncertainty to the best estimate

$$dp^{max} = \hat{dp}^{max} + \Delta dp^{max} = \left(1 + \frac{\Delta dp^{max}}{\hat{dp}^{max}}\right) \hat{dp}^{max} = \left(1 + \frac{\Delta dp^{max}}{\hat{dp}^{max}}\right) \frac{\hat{A}_d \hat{\beta}_d}{\hat{A}_h \hat{l}_{ds}} l^d. \tag{5.23}$$

Doing this will in theory always ensure that there is enough built-up pressure in the drill string to work with, at the expense of a smaller HeaveLock™ opening. However, the more certain the drill string and mud data, the less expense will lie in the result of this approach. A case example showing how data uncertainty plays a role in the result of this approach is found in Appendix A. The case example is based on very uncertain data, showing how a nominal differential pressure still can be computed with next to no knowledge of mud and drill string data. As an example, (1500 ± 500) kg/m^3 is a very uncertain measure for mud density, but it may still fit well with most drilling fluids that are used. In the case example the approach finds a maximum desired differential pressure with a relative uncertainty of 63.72%.

With this kind of measure for the desired differential pressure, though, we are at least able to find a nominal opening for HeaveLock™ with (5.9). However, this approach may be sufficient for vertical wells only, because that is the only case when the minimum length is a good estimate for the actual length, as the TVD will be close to the MD. For horizontal wells, this approach can only give a good estimate for the minimum length of the drill string, i.e. the vertical part, and will only suffice as a measure for the maximum differential pressure needed. As drilling continues in a horizontal plane, the drill string length will increase and the needed differential pressure would decrease, allowing for a larger valve opening. The case example (Appendix A) mentioned above also shows the difference this approach can have in a vertical and horizontal well. If a horizontal well is substantially longer in MD than TVD, this approach would find an unnecessary high differential pressure.

Counting Connections

Another approach to finding the length of the drill string is to implement an algorithm that counts the number of connections that are made. If all stands that are added to a drill string during drilling are of the same length, the count multiplied with the length and some elasticity constant will be close to a perfect measure of the drill string's length. In reality, however, not all connections will add to the drill string with the same length. Some joints could have been redone due to imperfections, the Bottom Hole Assembly (BHA) could have units or tools of varying length, or the driller might just add a single joint instead of a whole stand, for whatever reason. Although this uncertainty is an issue, it could very well outperform the also uncertain minimum length approach presented above, especially when drilling horizontal wells.

The robustness of such an algorithm is also very important. What happens if the algorithm loses its count somehow, e.g. if the on-board computer is faulty or reboots? Then, an approach that relies on physical measurements would be much more reliable.

Active Sonar Technology

Another potential solution could perhaps be to use active sonar technology to measure the length of the drill string. By emitting sound waves through the mud from a transmitter downhole, we can measure the time it takes for these sound waves to reflect back to a receiver downhole, see Fig. 5.8. If we know the speed of sound for the drilling fluid we can then calculate how far the signal has traveled, which must be equivalent to twice the length of the drill string [23].

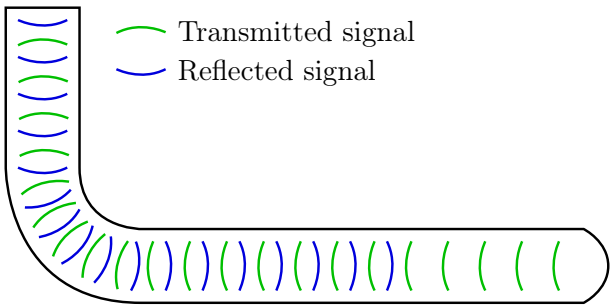

Figure 5.8: Active sonar principle in a horizontal well

I have not researched how signals will reflect in a drill string with bends and transitions between joints with different inner diameters. However, there is a communication technology called mud pulse telemetry that uses sound waves to transfer data from downhole to the topside and vice versa [24]. If the reflected signal from active sonar technology is found to be unreadable, mud

pulse telemetry could be a viable option to replicate this technology. In the simplest form, a downhole transmitter could transmit a message that a topside receiver can recognize, only for a topside transmitter to transfer the message back to a downhole receiver. If one were to implement such a technology, one could argue that this can be used to just transmit the actual length in the first place, as the length is probably already known to systems topside. The only downside would be that the system then relies on other existing systems.

5.3.3 Implementation

Whichever strategy that is chosen to estimate the length of the drill string, let us now just assume that the length is known within some certainty and that we are able to calculate the desired differential pressure because of it. We can then find a nominal opening for HeaveLock™ with (5.9) and eliminate drift by punishing deviation from the nominal value. However, our strategy to exploit the compressibility of fluids to calculate the desired differential pressure in Section 5.3.1 is actually independent of the mud-flow, while calculation of a nominal opening is not. It would be a better idea to punish deviation from the desired differential pressure instead, rather than the nominal opening, as the mud-flow could be subject to change, but the needed differential pressure would stay the same.

One way of punishing such deviations from nominal values is the use of an integrator in the control loop. If whatever value deviates from its desired value, an integrator term will add to (or subtract from, subject to the error's polarity) the control input until the value indeed reaches its intended value. In this case, where we would like to punish deviation from the desired differential pressure, we introduce the error variable

$$e_{\Delta p_{hl}} = \Delta p_{hl} - dp^d, \tag{5.24}$$

where $\Delta p = p_{hl0} - p_{hl}$. The error will be positive if we would like to reduce the differential pressure and vice versa. An integral term can be written as

$$y_{int} = \frac{K_p}{T_i s} e_{\Delta p_{hl}}, \tag{5.25}$$

where K_p is the proportional gain and T_i is the integral time. We multiply both the numerator

and denominator with $X(s)$ to get the two equations

$$y_{int} = K_p X(s) \tag{5.26a}$$

$$e_{\Delta p_{hl}} = T_i s X(s), \tag{5.26b}$$

and with inverse Laplace we get the differential and measurement equations

$$\dot{x} = \frac{1}{T_i} e_{\Delta p_{hl}} = \frac{1}{T_i}(\Delta p_{hl} - dp^d) \tag{5.27a}$$

$$y_{int} = K_p x, \tag{5.27b}$$

which also can be drawn as a block diagram, see Fig. 5.9. The block diagram also shows which inputs that are needed when we implement the integrator in the HeaveLock™ controller's state-space system.

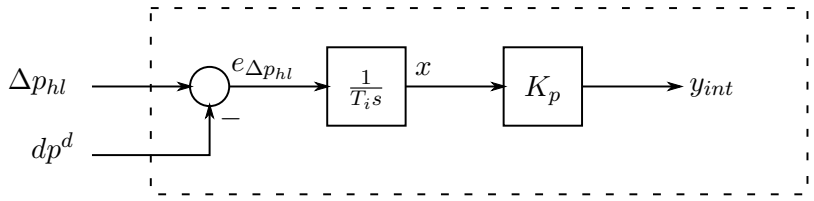

Figure 5.9: Integrator part of state-space system

The integrator in (5.27) can be added to the controller's state-space system (5.1a) by adding another state for the integrator and adding the integrator output to the y_2 term in (5.2)

$$\dot{x}_{n+1} = \frac{1}{T_i}\left((p_{hl0} - p_{hl}) - dp^d\right) = \boldsymbol{b}_u^\top \boldsymbol{u} - \frac{1}{T_i} dp^d \tag{5.28a}$$

$$y_2 = \boldsymbol{c}_2^\top \boldsymbol{x} + y_{int} = \boldsymbol{c}_2^\top \boldsymbol{x} + K_p x_{n+1}, \tag{5.28b}$$

where $\boldsymbol{b}_u^\top = \begin{bmatrix} \frac{1}{T_i} & -\frac{1}{T_i} & 0 & 0 & 0 & 0 \end{bmatrix}$, $\boldsymbol{c}_2^\top \boldsymbol{x}$ is the old y_2 output and $\boldsymbol{c}_2^\top$ is the second row of the old output matrix $\boldsymbol{C}$. The new state-space system is then written as

$$\dot{\boldsymbol{x}} = \begin{bmatrix} \boldsymbol{A} & \boldsymbol{0}_{n\text{x}1} \\ \boldsymbol{0}_{1\text{x}n} & 0 \end{bmatrix} \begin{bmatrix} \boldsymbol{x} \\ x_{n+1} \end{bmatrix} + \begin{bmatrix} \boldsymbol{B} & \boldsymbol{0}_{n\text{x}1} \\ \boldsymbol{b}_u^\top & b_r \end{bmatrix} \begin{bmatrix} \boldsymbol{u} \\ r \end{bmatrix} \tag{5.29a}$$

$$\boldsymbol{y} = \begin{bmatrix} \boldsymbol{c}_1^\top & 0 \\ \boldsymbol{c}_2^\top & K_p \end{bmatrix} \begin{bmatrix} \boldsymbol{x} \\ x_{n+1} \end{bmatrix} + \begin{bmatrix} \boldsymbol{D} & \boldsymbol{0}_{2\text{x}1} \end{bmatrix} \begin{bmatrix} \boldsymbol{u} \\ r \end{bmatrix}, \tag{5.29b}$$

where $b_r = -\frac{1}{T_i}$ and $r = dp^d$.

5.3.4 Anti-windup

Windup is the term used when an integrator keeps integrating when it does not need to, e.g. when it does not have any effect. This can happen if the controlled state does not reach its desired value fast enough. As an example, let us imagine a level controlled water tank with a controllable valve as input and a varying output. If we want to decrease the level of the tank, the controlled valve must close. Integral action will, in this case, try to decrease the valve opening until the water level has reached its desired height. Windup occurs when the valve is already fully closed, but the water level has yet to reach its desired height. Without any anti-windup measure, the integral term will keep increasing, trying to close the valve even further without any effect. The problem occurs when the water level reaches its desired height and the valve must open again to keep it there. The integral term may have grown so large that it will prevent the valve to open at first. An anti-windup measure is some kind of logic that will prevent this behavior by ceasing integration when a saturation limit has been met.

The flow through HeaveLock™ can never be negative, as the one-way valve in the drill string will prevent it. This is a saturation limit that may lead to windup problems, and an anti-windup measure should be considered. The idea is that if the desired flow is less than zero, there is no need for an integrator to further try to decrease the desired flow when the actual flow cannot be any less anyway.

Let us say that we have a persistent differential pressure error of $e_{\Delta p_{hl}}$. The integrator will then increase its output with $\frac{K_p}{T_i} e_{\Delta p_{hl}}$ every second. This is the very nature of an integrator. This means that the desired flow will change with the amount of $\frac{K_p}{T_i} e_{\Delta p_{hl}}$ every second. If now, after some time, the flow reaches its saturation limit, the integrator will still increase its output with $\frac{K_p}{T_i} e_{\Delta p_{hl}}$. There is no point in this, because the flow is already saturated, and the integrator should stop integrating. An easy way to do this is to add a value x to the integrator, i.e.

$$y_{int} = \frac{K_p}{T_i s} \left(e_{\Delta p_{hl}} + x \right), \qquad \text{where} \qquad x = \begin{cases} 0, & \text{if } q_{hl}^d = q^{sat} \\ -e_{\Delta p_{hl}}, & \text{otherwise,} \end{cases} \tag{5.30}$$

where q^{sat} is the saturated flow value. To fulfill the conditions for x, we can choose it as

$$x = K_{aw}(q^{sat} - q_{hl}^d) = -e_{\Delta p_{hl}}, \tag{5.31}$$

where we assume that the difference $q_{hl}^d - q^{sat}$ will be due to the integrator windup, i.e.

$$-\left(q^{sat} - q_{hl}^d\right) = \frac{K_p}{T_i} e_{\Delta p_{hl}}, \tag{5.32}$$

which means that the anti-windup gain can be chosen as

$$K_{aw} = \frac{T_i}{K_p}, \tag{5.33}$$

and we have that

$$y_{int} = \frac{K_p}{T_i s}\left(e_{\Delta p_{hl}} + \frac{T_i}{K_p}(q^{sat} - q_{hl}^d)\right), \tag{5.34}$$

which means that (5.27) can be changed to

$$\dot{x} = \frac{1}{T_i}(\Delta p_{hl} - dp^d) + \frac{1}{K_p}(q^{sat} - q_{hl}^d) \tag{5.35a}$$

$$y_{int} = K_p x. \tag{5.35b}$$

This can also be drawn as a block diagram as shown in Fig. 5.10. The block diagram shows how the integrator can be implemented in the controller's state-space system, as well as which inputs the system needs.

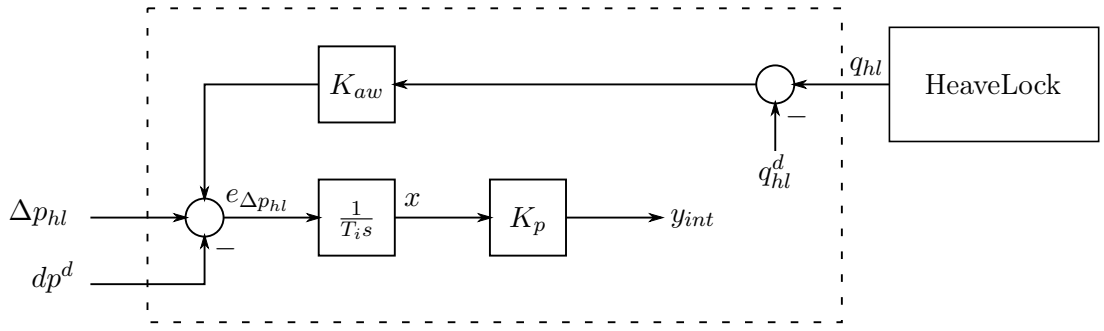

Figure 5.10: Integrator part of state-space system with anti-windup

If we have knowledge of the saturation limits, the saturated flow q^{sat} can just be computed directly in the integrator algorithm, e.g. as `max(0.0, q^d)` if negative flow is the only saturation. Or better yet, just freeze the integrator when you know that the saturation is met. However, if the saturation limits are not known, one should use the measured flow as the saturated flow, since the actual flow is what defines the saturation in the first place. If the dynamics of the system is slow, though, this could worsen the performance of the integrator, as it would take time for the actual value to reach the desired value, making $q_{hl}^d \neq q^{sat}$ even though the flow is not saturated. The integral time and anti-windup gain should then be adjusted for optimal performance. If the

measurement is used, the HeaveLock™ controller's state-space system in (5.29) can be changed to

$$\dot{\boldsymbol{x}} = \begin{bmatrix} \boldsymbol{A} & \boldsymbol{0}_{n\text{x}1} \\ \boldsymbol{0}_{1\text{x}n} & 0 \end{bmatrix} \begin{bmatrix} \boldsymbol{x} \\ x_{n+1} \end{bmatrix} + \begin{bmatrix} \boldsymbol{B} & \boldsymbol{0}_{n\text{x}2} \\ \boldsymbol{b}_u^\top & \boldsymbol{b}_r^\top \end{bmatrix} \begin{bmatrix} \boldsymbol{u} \\ \boldsymbol{r} \end{bmatrix} \tag{5.36a}$$

$$\boldsymbol{y} = \begin{bmatrix} \boldsymbol{c}_1^\top & 0 \\ \boldsymbol{c}_2^\top & K_p \end{bmatrix} \begin{bmatrix} \boldsymbol{x} \\ x_{n+1} \end{bmatrix} + \begin{bmatrix} \boldsymbol{D} & \boldsymbol{0}_{2\text{x}2} \end{bmatrix} \begin{bmatrix} \boldsymbol{u} \\ \boldsymbol{r} \end{bmatrix}, \tag{5.36b}$$

where $\boldsymbol{b}_u^\top = \begin{bmatrix} \frac{1}{T_i} & -\frac{1}{T_i} & 0 & \frac{1}{K_p} & 0 & 0 \end{bmatrix}$, $\boldsymbol{b}_r^\top = \begin{bmatrix} -\frac{1}{T_i} & -\frac{1}{K_p} \end{bmatrix}$ and $\boldsymbol{r} = \begin{bmatrix} dp^d & q_{hl}^d \end{bmatrix}$.

Fig. 5.11 shows the simulation results from a connection where an integrator with anti-windup was implemented in HeaveLock™'s controller. Equation (5.13), together with the minimum length approach in Section 5.3.2, was used to find the desired differential pressure using a desired length of $l^d = 5$ meters. Compared to Fig. 5.5 and 5.6 we see that the integrator is able to eliminate drift in differential pressure and choke opening.

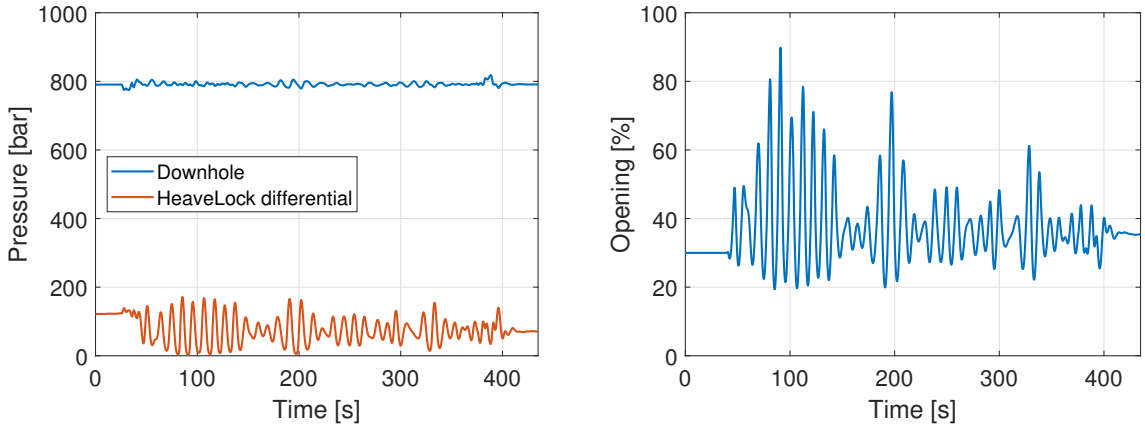

(a) Pressures (b) HeaveLock™'s opening

Figure 5.11: Integrator performance with desired length of 5 meters

The heave data used in the simulation was the same as the one we analyzed in Section 4.1. We see from Fig. 4.1 that the maximum peak-to-peak value of heave is only about 2.5 meters, only half of what was used to produce the results presented in Fig. 5.11. Fig. 5.12 shows the same simulation, only now with a desired length of $l^d = 2.5$ meters. This decreases the desired differential pressure by half, and as expected we see a greater nominal opening because of it in Fig. 5.12b. However, matching the desired length with the actual maximum peak-to-peak amplitude of the heave motion drives HeaveLock™ to saturation, losing differential pressure over HeaveLock™ and in effect decreasing its performance. This means that the desired length must be somewhat greater than the peak-to-peak value of the actual heave.

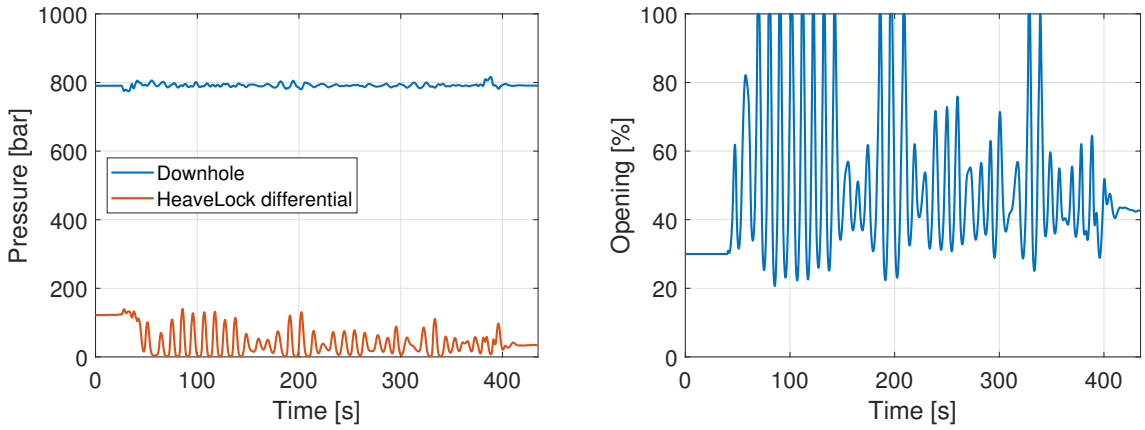

(a) Pressures (b) HeaveLock™'s opening

Figure 5.12: Integrator performance with desired length of 2.5 meters

5.4 Dynamic Correction

The drift solution presented in Section 5.3.1 relies on having an idea about the desired length l^d of axial motion that HeaveLock™ should be able to control, e.g. the maximum allowable peak-to-peak value of the heaving motion for drilling to continue. Equation (5.13) will then find a maximum desired differential pressure over HeaveLock™ based on this knowledge. However, most of the time, the actual peak-to-peak value of the heaving motion will be less than the maximum allowable value, meaning that the desired differential pressure could be less.

By integrating the already available/measured bit velocity we can find the bit displacement. The bit displacement gives us information about the length the bit is moving axially along the well. If this data is logged, we can use the historical data and statistics to gain valuable information about the bit motion. A common way to keep track of live data, without the need for extreme amounts of memory, is using a moving window algorithm. A moving window algorithm always keeps track of the latest applicable data, see Fig. 5.13. The figure illustrates historical bit motion data to the left of $t = 0$ and future data to the right. The moving window contains the latest of the historical data, where the amount of data is determined by a predefined window width. Such an algorithm is often used to calculate so-called moving averages, where the moving average, of course, is the average of the data within the moving window. Just like any other numerical data, though, the data contained in the moving window can be used to so much more [25].

If we calculate the average and standard deviation of historical data, it can be used as a measure for the expected data in the near future. A moving standard deviation will give us a dynamic measure with the ability to adapt to changes, e.g. developing weather conditions. By definition,

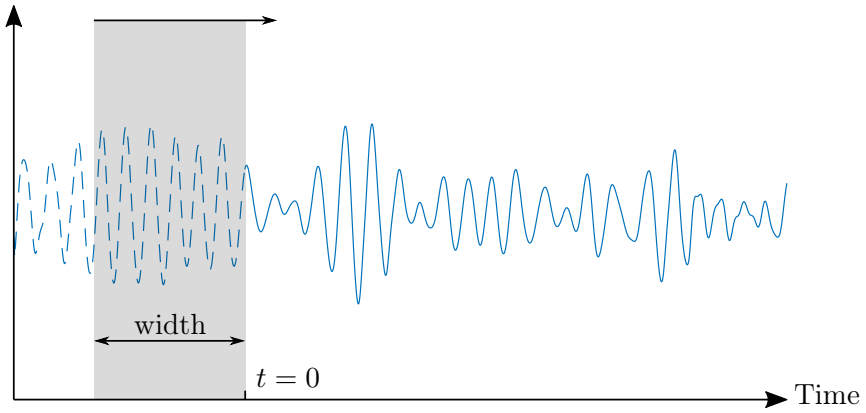

Figure 5.13: Moving window algorithm

if the gathered data is approximately normally distributed, about 68% of the data will lie within one standard deviation σ of the mean μ, and over 99% within 3 standard deviations [26]. This means that 6σ should give a good estimate of the maximum peak-to-peak amplitude of the gathered data. The heaving motion is not necessarily normally distributed, though, but the peak-to-peak amplitude can still be defined by the standard deviation, i.e.

$$A_{pp} = c_{pp}\sigma, \tag{5.37}$$

where c_{pp} is some constant value. Fig. 5.14 shows an example of how a moving standard deviation algorithm can be used to find an estimate of the maximum expected peak-to-peak amplitude. Here, a threshold of $\pm 2.5\sigma$ is used to plot the expectancy interval, i.e. the estimated maximum peak-to-peak amplitude is $A_{pp} = 5\sigma$. The moving window used to produce the expectancy interval had a width of 100 seconds.

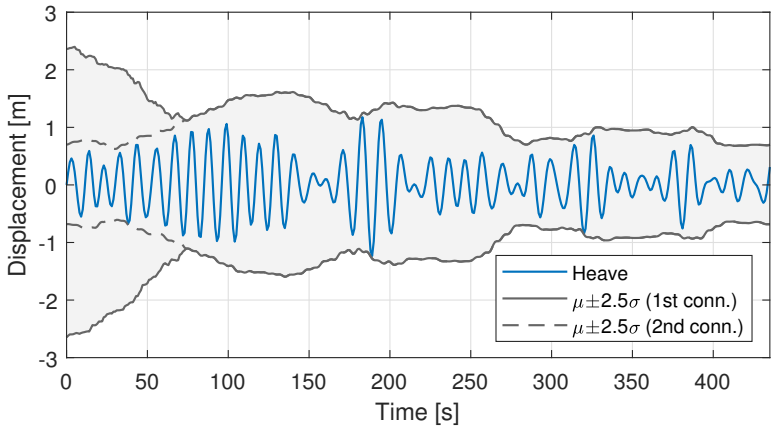

Figure 5.14: Moving standard deviation with a moving window width of 100 seconds

Fig. 5.14 also shows two different approaches to initialization of the moving window algorithm, labeled 1st and 2nd connection. For the first connection or the first time the algorithm is activated, it is difficult to say something about the expected values since the window with historical data is empty. If one were to initialize the window with zeroes, the expected mean and standard deviation would also be zero, resulting in a peak-to-peak amplitude, and thus also the desired differential pressure, of zero. The desired opening of HeaveLock™ at initialization would then be fully open, and HeaveLock™ would not be able to compensate for the first occurring downhole pressure oscillations. The window should rather be initialized with data that represents the maximum allowed bit motion, such that there is enough built-up pressure to work with when HeaveLock™ is activated. The standard deviation will then converge down to its actual value rather than up, see the example of the 1st connection in the figure. The next time HeaveLock™ is activated, i.e. for the 2nd connection, the window already contains data from the previous connection. If the time between connections is low, it should be safe to assume that the weather has not changed dramatically and that data from the previous connection is a good estimate for the next. However, if the time between connections is high for whatever reason, the initialization of the 1st connection is the safe approach.

The desired length that HeaveLock™ should be able to control can be defined as

$$l^d = \alpha\sigma = K_l A_{pp}, \tag{5.38}$$

where $\alpha = K_l c_{pp}$ and K_l is a design parameter. Together with (5.13) we now have a dynamic desired differential pressure that is a function of the standard deviation, rather than an uncertain constant, i.e.

$$dp^d = \frac{A_d \beta_d}{A_h l_{ds}} \alpha\sigma. \tag{5.39}$$

The results of a simulation using the moving standard deviation algorithm described above is shown in Fig. 5.15. The estimated maximum peak-to-peak amplitude was chosen as $A_{pp} = 6\sigma$, and the desired length was chosen as $l^d = 2A_{pp} = 12\sigma$ in the simulation. HeaveLock™ was activated at $t = 40$ and deactivated at $t = 375$ seconds.

As we can see in Fig. 5.15a, the desired length is initiated fairly high, but is able to adapt to the logged bit displacement data over time. The logged bit motion never comes close to half the desired length, though, and one could argue that K_l should be decreased. However, if we look at Fig. 5.15d, we see that HeaveLock™ actually saturates around $t = 330$ seconds. We also see from Fig. 5.15b that the differential pressure is lost when this happens. This is not necessarily

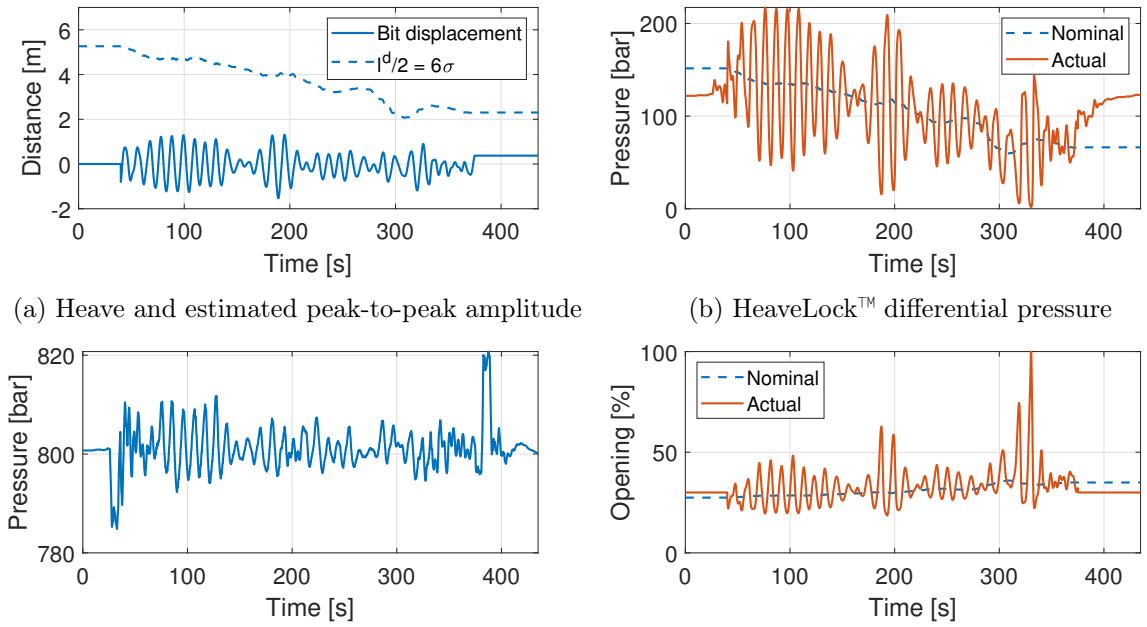

(a) Heave and estimated peak-to-peak amplitude

(b) HeaveLock™ differential pressure

(c) Downhole pressure

(d) Choke opening

Figure 5.15: Simulation with dynamic estimation of peak-to-peak amplitude

an issue if it happens for only a short amount of time when the bit motion is changing direction, but if it were to persist the performance of HeaveLock™ would be decreased. This means that we cannot afford to decrease K_l, as this would only aggravate the issue. Putting this aside, we can see that the strategy works as intended. The moving standard deviation algorithm adapts the desired length to fit the actual bit motion, giving us a dynamic nominal differential pressure and opening for HeaveLock™. With the proper design parameters, this strategy should give us an optimal differential pressure that minimizes wear and tear and pump pressure while HeaveLock™ is active.

Note that the act of hoisting and lowering of the drill string will influence the moving standard deviation significantly, both increasing the desired differential pressure and decreasing HeaveLock™'s opening. In Fig. 5.15 HeaveLock™ was activated after the drill string was hoisted and deactivated before it was lowered back down to prevent this. Another solution would be to use the band-pass filter defined in (2.10) to attenuate the low-frequency hoisting/lowering as well as any high-frequency noise. Fig. 5.16 shows the effect that lowering of the drill string has if HeaveLock™ stays activated, both with and without the mentioned band-pass filter implemented. A bandwidth w_{bw} of 0.6 and center frequency w_c of 0.63 rad/s was used to produce the results. Without the band-pass filter, the standard deviation increases as the drill string is lowered, increasing the desired differential pressure up until the point where HeaveLock™'s opening is fully closed. The band-pass filter prevents this with next to no impact on performance.

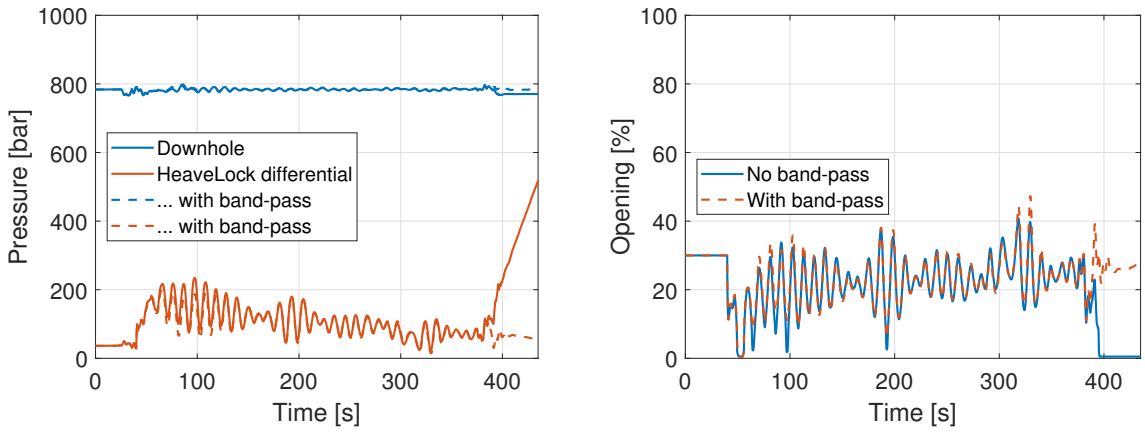

(a) Pressures (b) HeaveLock™'s opening

Figure 5.16: The advantage of having a band-pass filter while hoisting/lowering

6 Joint Operation

Now that we have a controller for HeaveLock™ that we can imagine being implemented in a real application, i.e. it reduces the heave-induced downhole pressure oscillations during connections and drift is eliminated, we can begin examining the effect of HeaveLock™ on the MPD system and vice versa. As HeaveLock™ only will be active while performing connections and preferably be fully open otherwise, the only interesting response to study the performance of is connections. Simulations of connections both with and without the MPD system and HeaveLock™ implemented are presented and analyzed in the following section. The response with both the static desired length strategy in Section 5.3.1 and the dynamic correction strategy in Section 5.4 will also be compared. Hopefully, we will not see any changes in neither the MPD system's nor HeaveLock™'s performance during joint operation.

6.1 Performance

Fig. 6.1 shows a comparison of the simulation results from connections done with and without both the MPD system presented in Section 3 and HeaveLock™ implemented. Both a static nominal opening and the dynamic correction strategy presented in Section 5.4 was used in HeaveLock™'s controller in the simulations, and the results are shown in Fig. 6.1a and 6.1b, respectively, and perhaps easiest seen from HeaveLock™'s nominal opening in Fig. 6.1a3 and 6.1b3.

If we compare HeaveLock™'s opening with and without MPD, they are almost identical. This is true for both the static and dynamic nominal opening case, and is not surprising, as the implemented controller in HeaveLock™ is purely based on the bit velocity measurement. As long as the heaving motion is identical, the opening should, therefore, stay the same. This means that HeaveLock™ should be able to operate independently of the MPD system, as long as its controller is only dependent on the bit velocity. The MPD system's controller, however, is dependent of states that HeaveLock™ may affect, such as the choke pressure and choke flow, as seen from (3.2) and (3.3). As we know, HeaveLock™ is designed to reduce the oscillations in downhole pressure, which means that the pressure waves transferred to the topside choke should also be reduced. This should, in theory, result in less variation in choke pressure error and thus desired choke flow, as well as a calmer choke opening, according to (3.2) and (3.3).

We also notice that HeaveLock™ is able to decrease the heave-induced downhole pressure oscil-

lations both with and without the MPD system implemented, no matter if the nominal opening is static or dynamic. The decreased oscillations are also very similar regardless of whether MPD is implemented or not, a good indication that joint operation should not present any big problems.

Fig. 6.1a2 and 6.1b2 shows the topside choke opening during the simulated connection. With a static nominal opening, it is difficult to spot much difference in the opening, expect just after HeaveLock™ is activated. The temporary greater topside choke's opening from around $t = 50$ s in Fig. 6.1a2 comes from HeaveLock™'s nominal opening being greater than its initial opening at activation. HeaveLock™ will start with increasing its opening and end up operating around a higher nominal value. This will decrease the differential pressure over HeaveLock™, resulting in a temporary increased flow due to the volume inside the drill string expanding. The increased flow will, in turn, increase the differential pressure over the topside choke, meaning it must increase its opening to compensate for the pressure change. The opposite is evident in Fig. 6.1b, where HeaveLock™'s nominal opening is smaller than its initial opening at activation. This time HeaveLock™ starts with decreasing its opening, temporarily decreasing the flow, and the topside choke reacts to the decreased flow by closing its opening at around $t = 50$ s.

The bump in the topside choke's opening from $t = 300$ s (Fig. 6.1b2) can also be explained with the same reasoning. Recent low amplitude heave motion led to a temporary decreased moving standard deviation in the dynamic correction algorithm (see the bump in nominal opening in Fig. 6.1b3 at the same time), temporarily increasing the flow and thus the topside choke's opening. This effect can actually be reduced by extending the moving window in the moving standard deviation algorithm presented in Section 5.4.

If we look closely, we also see that the topside choke opening seems to operate around a slightly higher opening than what is the case of having a static nominal opening. This is the result of the nominal opening of HeaveLock™ slightly increasing throughout the whole time of the connection, in effect having a constantly higher flow than when using a static nominal opening due to the expanding volume.

From Fig. 6.1 it seems like the only state affecting the performance of the MPD system is the nominal flow determined by manipulation of the nominal opening of HeaveLock™. Having HeaveLock™ operate around a static nominal opening does not really change the performance at all since the nominal flow stays the same. Small constant changes in HeaveLock™'s nominal opening only increases the nominal flow by some constant, and the topside choke simply compensates for

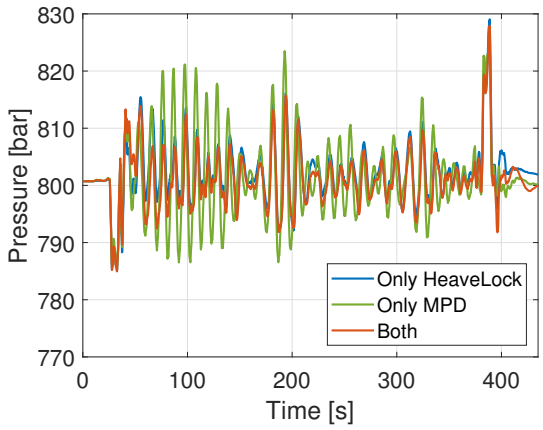

(a1) Downhole pressure

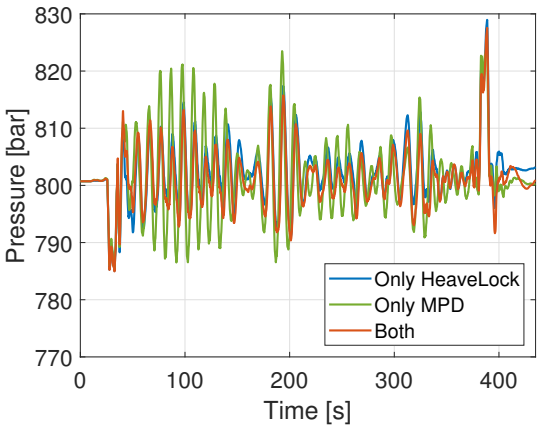

(b1) Downhole pressure

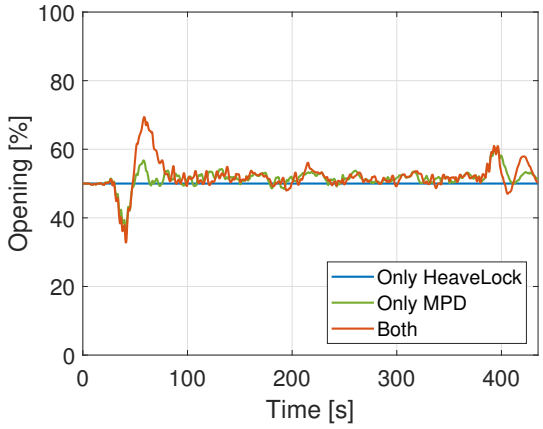

(a2) Topside choke opening

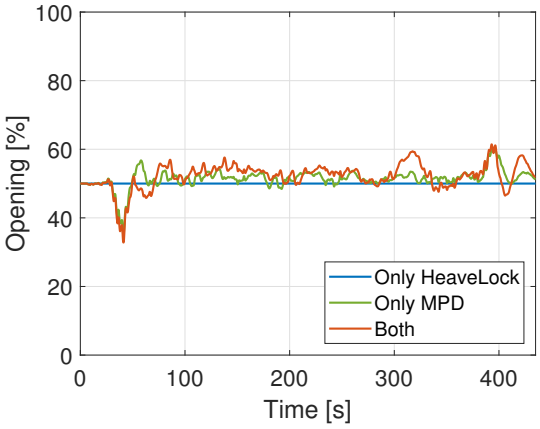

(b2) Topside choke opening

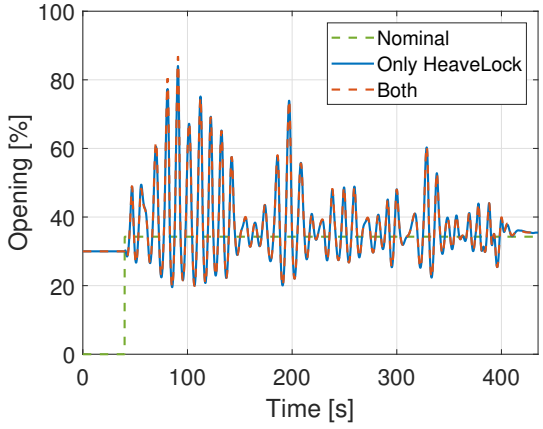

(a3) HeaveLock™ opening

(a) Static nominal opening

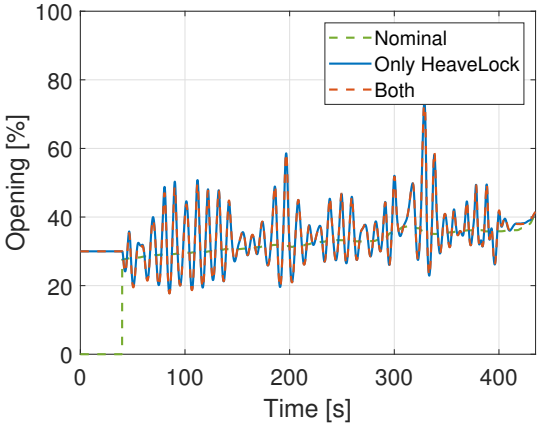

(b3) HeaveLock™ opening

(b) Dynamic nominal opening

Figure 6.1: Joint operation of MPD and HeaveLock™ both with and without dynamic correction

this by finding a new nominal value to operate around. Sudden changes in HeaveLock™'s nominal opening is what influence the MPD system's response the most, but not necessarily worsening its performance. Without the topside choke reacting to persistent flow changes, the back-pressure and thus downhole pressure should increase or decrease, depending on the flow change being positive or negative, respectively. The MPD system actually ensures that the downhole oscillations are kept around the desired downhole pressure by compensating for any persistent flow changes created by HeaveLock™. This is somewhat evident around $t = 60$ s in Fig. 6.1a1 and $t = 320$ s in Fig. 6.1b1, where the downhole pressure is a little lower with both systems implemented, than with only HeaveLock™.

The effect that the nominal flow has on the MPD system is most noticeable at the activation of HeaveLock™, because of the sudden change in the nominal opening. If sudden changes are big enough, one could find that the topside choke saturates trying to compensate for the sudden changing flow, but then again, without the MPD system the downhole pressure would change when flow changes are apparent, meaning that having a topside choke that saturates may still be better than having an opening that is not manipulated at all. To be on the safe side, HeaveLock™'s opening should decrease gradually to its initial nominal opening at activation, as well as increase gradually back to whatever opening is best when HeaveLock™ is not active. The filter that was suggested as a modification of the MPD system in Section 4 may even be modified further to attenuate the effects of ramping by adding a high-pass filter.

Other than that, the simulation results do not give rise to any suspicion that the joint operation of MPD and HeaveLock™ cannot be done, not with the control strategies mentioned in this thesis anyway. If HeaveLock™'s controller were to use measurements of a state affected by the MPD system rather than the bit velocity, e.g. of the downhole pressure, the results may have been different, requiring further analysis.

7 Concluding Remarks

7.1 Conclusion

Joint operation of MPD systems and HeaveLock™ on floating drilling rigs is possible, and most probably without the need for much coordination. MPD systems must, however, have functionality that prevents them to try compensating for pressure variations in the frequency spectrum of waves. We have seen that, without such functionality, MPD systems may aggravate the problem with heave-induced surge and swab pressures during connections, and found that the frequencies of heave must be attenuated. In this thesis, we have proposed a modification of a specific MPD system to prevent such behavior and seen from simulation data that the modified system does, indeed, no longer try to compensate for surge and swab. The suggested modification is an implementation of two separate filters to reduce the frequencies of heave in the choke valve's input. This made the choke valve's opening more or less stationary during connections, without significantly affecting the system's response to setpoint changes or pump flow disturbances. Implementing filters also had positive effects on the downhole pressure variations.

HeaveLock™'s original controller had a problem that caused its valve opening to drift and was also in need of modification. We have in this thesis developed a strategy to eliminate drift by exploiting physical knowledge about the system in addition to downhole measurements. Based on this knowledge we were able to calculate a nominal opening for HeaveLock™. The drift problem was then eliminated by punishing any deviation from this value by adding an integrator in the control loop. Next, we discussed that weather changes will affect the needed nominal opening and that a dynamic nominal opening based on live measurements would be preferred. We then came up with a strategy to dynamically calculate the needed nominal opening based on bit velocity measurements. The strategy proved successful in simulations. HeaveLock™ was able to compensate for surge and swab pressures while dynamically adapting to the changing heave conditions.

Lastly, we looked at the joint operation of the modified MPD and HeaveLock™ systems. From simulations, we saw that the MPD system did nothing to the response of HeaveLock™, simply because its controller is only dependent of the bit velocity and not any of the states affected by the MPD system. HeaveLock™ did, however, change the response of the MPD system slightly. The MPD system will try to compensate for any persistent flow changes due to the manipulation of HeaveLock™. The only time when this may be an issue is during ramping of HeaveLock™'s opening, e.g. at activation when pressure must be built up, or when returning to normal operations. All in all, joint operation of MPD and HeaveLock™ looks very promising.

7.2 Future Work

There was not spent much time researching how downhole equipment can find the length of the drill string on their own, for HeaveLock™ to be fully autonomous with regards to the strategy presented in Section 5.3.1, without receiving any information from topside equipment. If communicating such information is not acceptable, this topic should be addressed in later studies.

Another topic that was just briefly touched upon is the ramping of HeaveLock™'s opening at activation and before deactivation, as well as its effect on the MPD control system. How ramping should be performed is, to my knowledge, yet to be discovered, but once this is ascertained it is encouraged to study the effects that ramping will have on joint operation of the control strategies proposed in this thesis. If it turns out to be an issue for the MPD control system, one idea could be to modify the suggested filter such that disturbances due to ramping also will be attenuated.

Appendix A Case Example: Minimum Length

Let us say that at a depth of 4000 meters the pressure is measured to be 780 bar, we know that the operating window for back-pressure is maximum 20 bar and that the friction loss never exceeds 30 bar at this depth. Let's also assume that the mud weight can be within the interval from 1000 to 2000 kg/m^3, and a gravity that varies with ± 0.01 from 9.81 m/s^2.

$$\begin{aligned} p_{h,a} &= \hat{p}_{h,a} \pm \Delta p_{h,a} = \left(780 - \frac{1}{2}(20+30) \pm \frac{1}{2}(20+30)\right) \text{ bar} \\ \rho_a &= \hat{\rho}_a \pm \Delta\rho_a = (1500 \pm 500) \text{ kg/m}^3 \\ g &= \hat{g} \pm \Delta g = (9.81 \pm 0.01) \text{ m/s}^2. \end{aligned}$$

Let us assume that the relative errors defined by these numbers are good enough at any depth. I.e. the relative error for the minimum length is

$$\frac{\Delta l_{ds}^{min}}{l_{ds}^{min}} = \frac{\Delta p_{h,a}}{\hat{p}_{h,a}} + \frac{\Delta\rho_a}{\hat{\rho}_a} + \frac{\Delta g}{\hat{g}} = \frac{25}{755} + \frac{500}{1500} + \frac{0.01}{9.81} = 0.3675 = 36.75\%.$$

The best estimate for the minimum length is

$$\hat{l}_{ds}^{min} = \frac{\hat{p}_{h,a}}{\hat{\rho}_a \hat{g}} = \frac{755 \cdot 10^5}{1500 \cdot 9.81} = 5131 \text{ meters},$$

with an uncertainty of $\pm 36.75\%$, i.e. (5131 ± 1886) meters. In this case, the density is actually closer to 2000 kg/m^3, still within the certainty limits, i.e. the real length is closer to

$$\hat{l}_{ds}^{min} = \frac{\hat{p}_{h,a}}{\hat{\rho}_a \hat{g}} = \frac{755 \cdot 10^5}{1910 \cdot 9.81} = 4029 \text{ meters}.$$

Nonetheless, let us now assume that the bulk modulus is $(2.15 \pm 0.15) \cdot 10^5$ bar, the hydraulic area is 0.0093 m^2, the displaced area 0.0034 m^2, and both areas have an uncertainty of 10%. The uncertainty in maximum differential pressure is then

$$\frac{\Delta dp^{max}}{\hat{dp}^{max}} = \frac{\Delta A_d}{\hat{A}_d} + \frac{\Delta\beta_d}{\hat{\beta}_d} + \frac{\Delta A_h}{\hat{A}_h} + \frac{\Delta l_{ds}^{min}}{\hat{l}_{ds}^{min}} = 0.1 + \frac{0.15}{2.15} + 0.1 + 0.3675 = 0.6373 = 63.73\%,$$

This means that for any given measurement, the desired maximum differential pressure over HeaveLock™ should be at least 63.73% higher than the best estimate. If the peak-to-peak value of the heaving motion, and thus the desired length l^d, is 5.0 meters we then have a desired

maximum differential pressure of

$$dp^{max} = \left(1 + \frac{\Delta dp^{max}}{\hat{dp}^{max}}\right) \frac{\hat{A}_d \hat{\beta}}{\hat{A}_h \hat{l}_{ds}} l^d = 1.6373 \cdot \frac{0.0034 \cdot 2.15 \cdot 10^9}{0.0093 \cdot 5131} \cdot 5.0 = 12.54 \text{ bar.}$$

If the best estimate for the minimum length is lower, let us say about 1000 meters, the desired maximum differential pressure would be

$$dp^{max} = \left(1 + \frac{\Delta dp^{max}}{\hat{dp}^{max}}\right) \frac{\hat{A}_d \hat{\beta}}{\hat{A}_h \hat{l}_{ds}} l^d = 1.6373 \cdot \frac{0.0034 \cdot 2.15 \cdot 10^9}{0.0093 \cdot 1000} \cdot 5.0 = 64.35 \text{ bar.}$$

This also means that if the well in question was 5131 meters in MD, but only 1000 meters TVD (i.e. not vertical), this approach would set the desired maximum differential pressure a substantial amount higher than necessary, and thus also an unnecessary small HeaveLock™ opening.

Appendix B Discretization

B.1 Methods

Both Euler's forward and backward differentiation methods are simple approximations of the time derivative of a time-varying function.

Forward Euler

Euler's forward differentiation method is defined as

$$\dot{x} \approx \frac{x_{k+1} - x_k}{h}, \tag{B.1}$$

where h is the sampling period or step size equal to the time between steps $\Delta t = t_{k+1} - t_k$. It is mostly used for simulations, as it approximates the next state based on known information about the current state.

Backward Euler

Euler's backward differentiation method is defined as

$$\dot{x} \approx \frac{x_k - x_{k-1}}{h}, \tag{B.2}$$

where h is the sampling period or step size equal to the time between steps $\Delta t = t_k - t_{k-1}$. It is mostly used for discretizing signal filters and controllers, as it gives a current state based on known information about the previous state.

B.2 State-Space Models

A state-space representation of a model

$$\dot{\boldsymbol{x}} = \boldsymbol{A}\boldsymbol{x} + \boldsymbol{B}\boldsymbol{u} \tag{B.3a}$$

$$\boldsymbol{y} = \boldsymbol{C}\boldsymbol{x} + \boldsymbol{D}\boldsymbol{u} \tag{B.3b}$$

where $\boldsymbol{A}$, $\boldsymbol{B}$, $\boldsymbol{C}$ and $\boldsymbol{D}$ are matrices and $\boldsymbol{x}$, $\boldsymbol{u}$ and $\boldsymbol{y}$ are vectors. A state-space model is often used for simulations and discretized with forward Euler (B.1), i.e.

$$\frac{\boldsymbol{x}_{k+1} - \boldsymbol{x}_k}{h} = \boldsymbol{A}\boldsymbol{x}_k + \boldsymbol{B}\boldsymbol{u}_k \tag{B.4a}$$

$$\boldsymbol{x}_{k+1} = \boldsymbol{x}_k + h(\boldsymbol{A}\boldsymbol{x}_k + \boldsymbol{B}\boldsymbol{u}_k) \tag{B.4b}$$

$$\boldsymbol{x}_{k+1} = (\boldsymbol{I} + h\boldsymbol{A})\boldsymbol{x}_k + h\boldsymbol{B}\boldsymbol{u}_k, \tag{B.4c}$$

which gives the discrete state-space model

$$\boldsymbol{x}_{k+1} = \boldsymbol{A}_d\boldsymbol{x}_k + \boldsymbol{B}_d\boldsymbol{u}_k \tag{B.5a}$$

$$\boldsymbol{y}_k = \boldsymbol{C}_d\boldsymbol{x}_k + \boldsymbol{D}_d\boldsymbol{u}_k, \tag{B.5b}$$

where $\boldsymbol{A}_d = (1 + h\boldsymbol{A})$ and $\boldsymbol{B}_d = h\boldsymbol{B}$ are the discretized system and input matrices, respectively, $\boldsymbol{C}_d = \boldsymbol{C}$ is the output matrix and $\boldsymbol{D}_d = \boldsymbol{D}$ is the feed-forward matrix.

Appendix C MPD Performance Plots

List of Figures

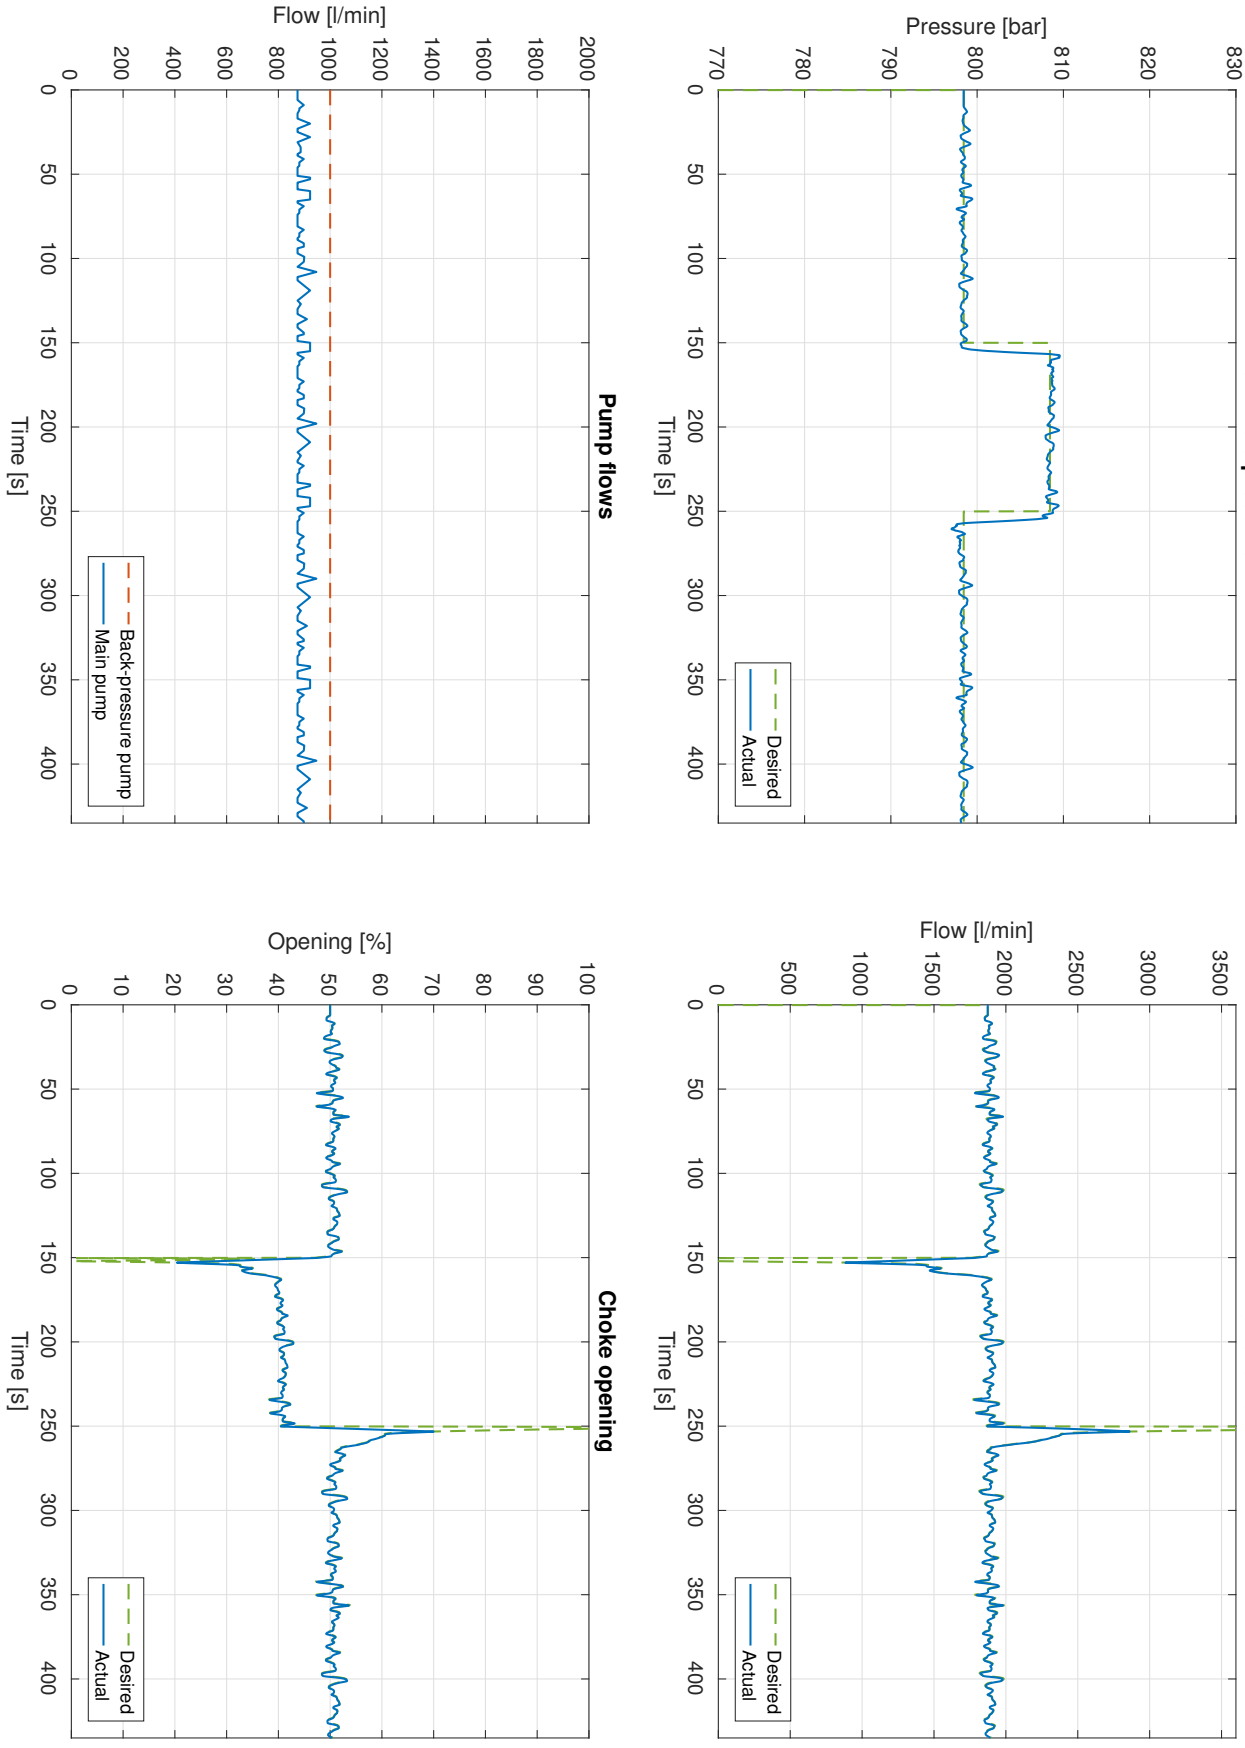

Figure C.1: Setpoint step response

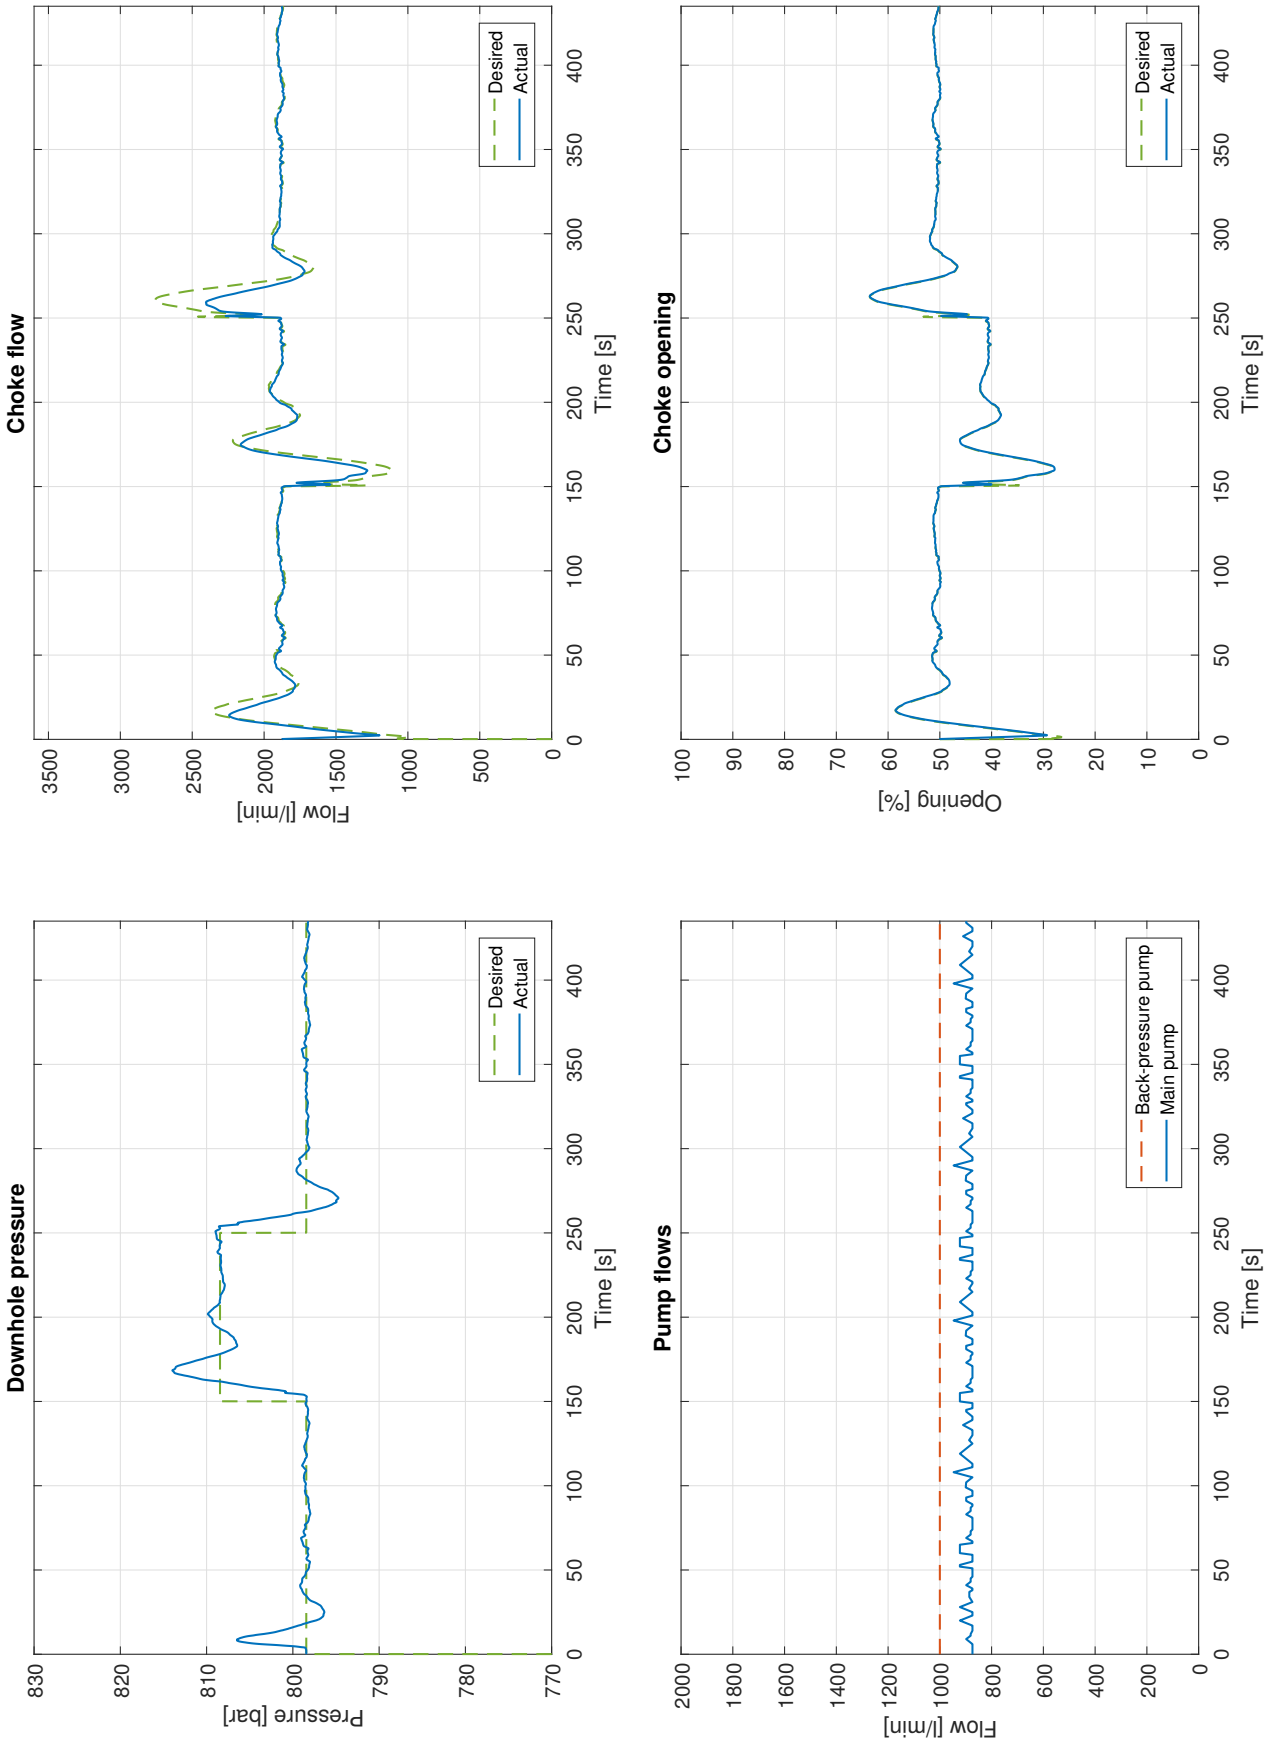

Figure C.2: Setpoint step response (filtered reference)

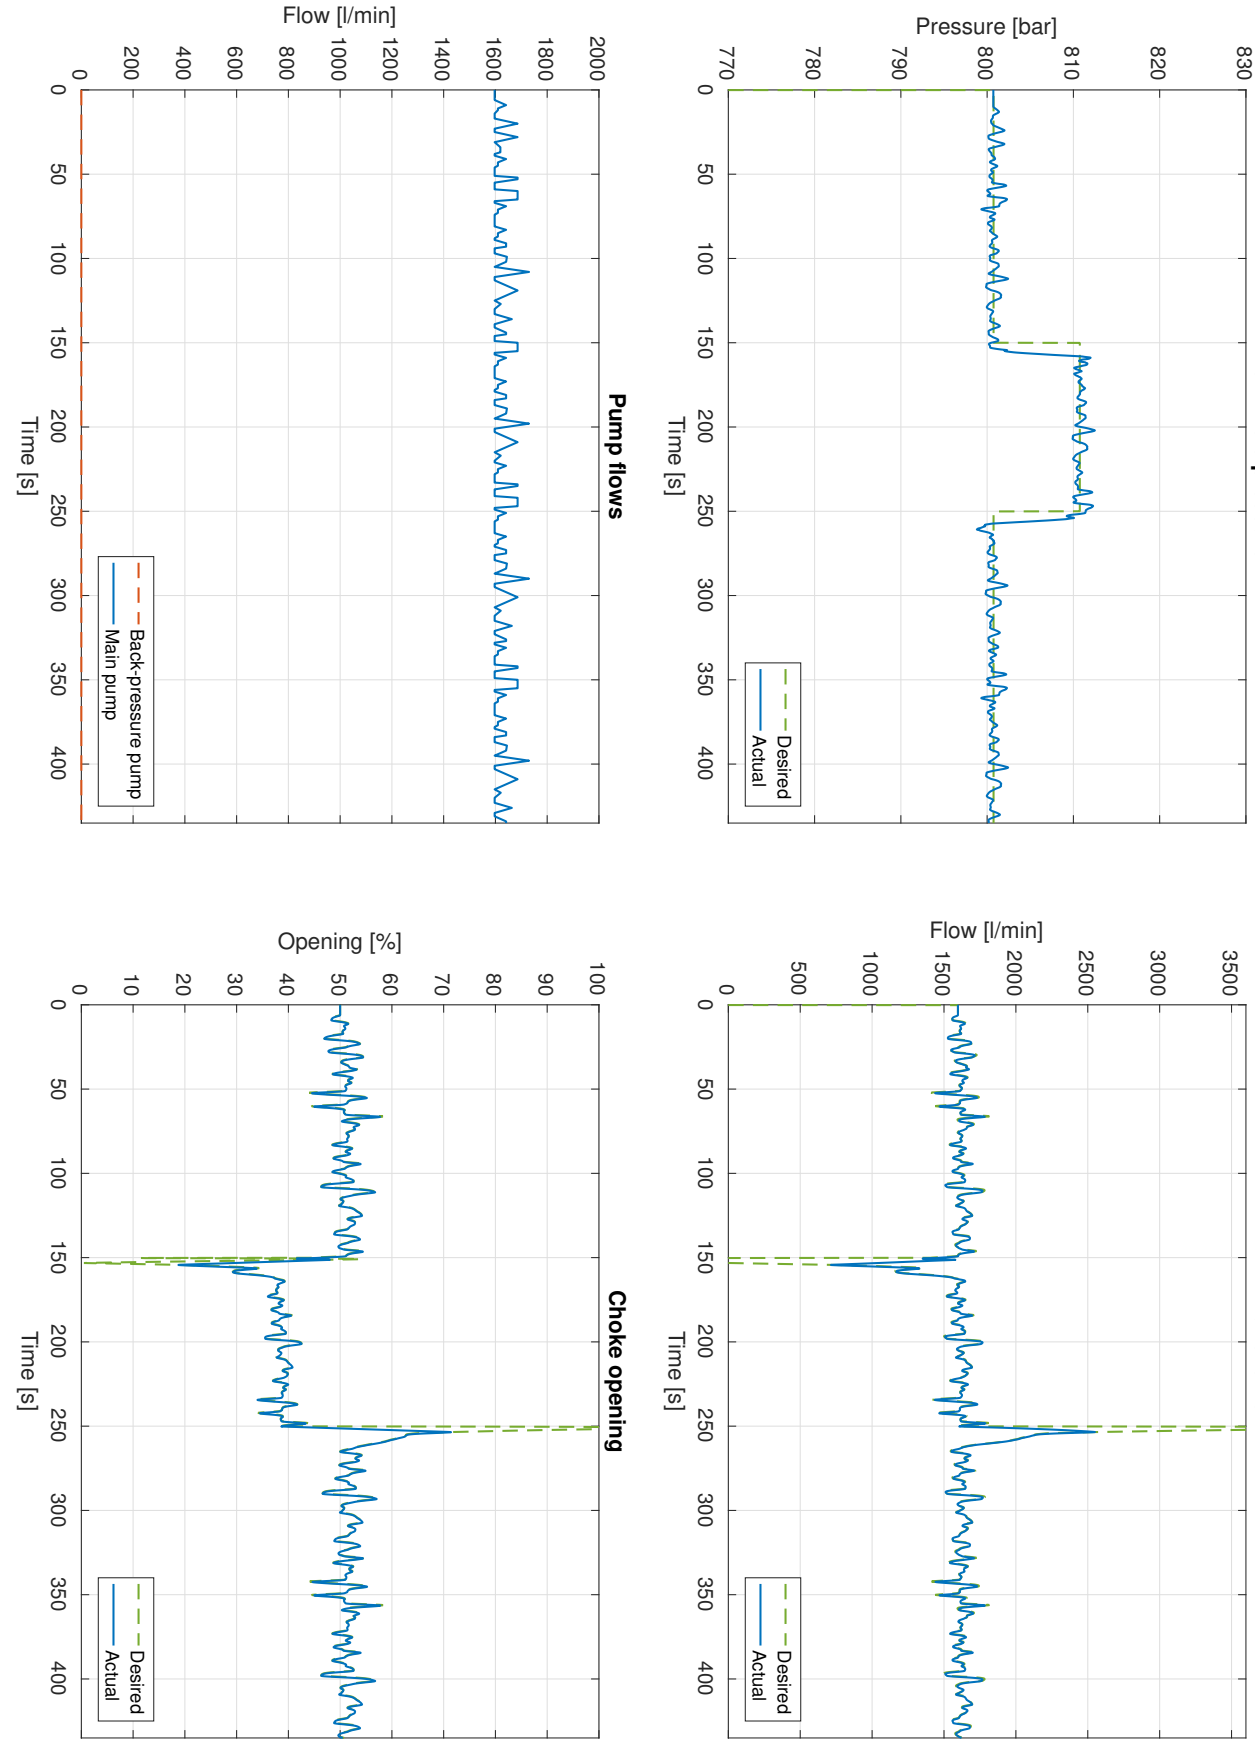

Figure C.3: Setpoint step response with CCS

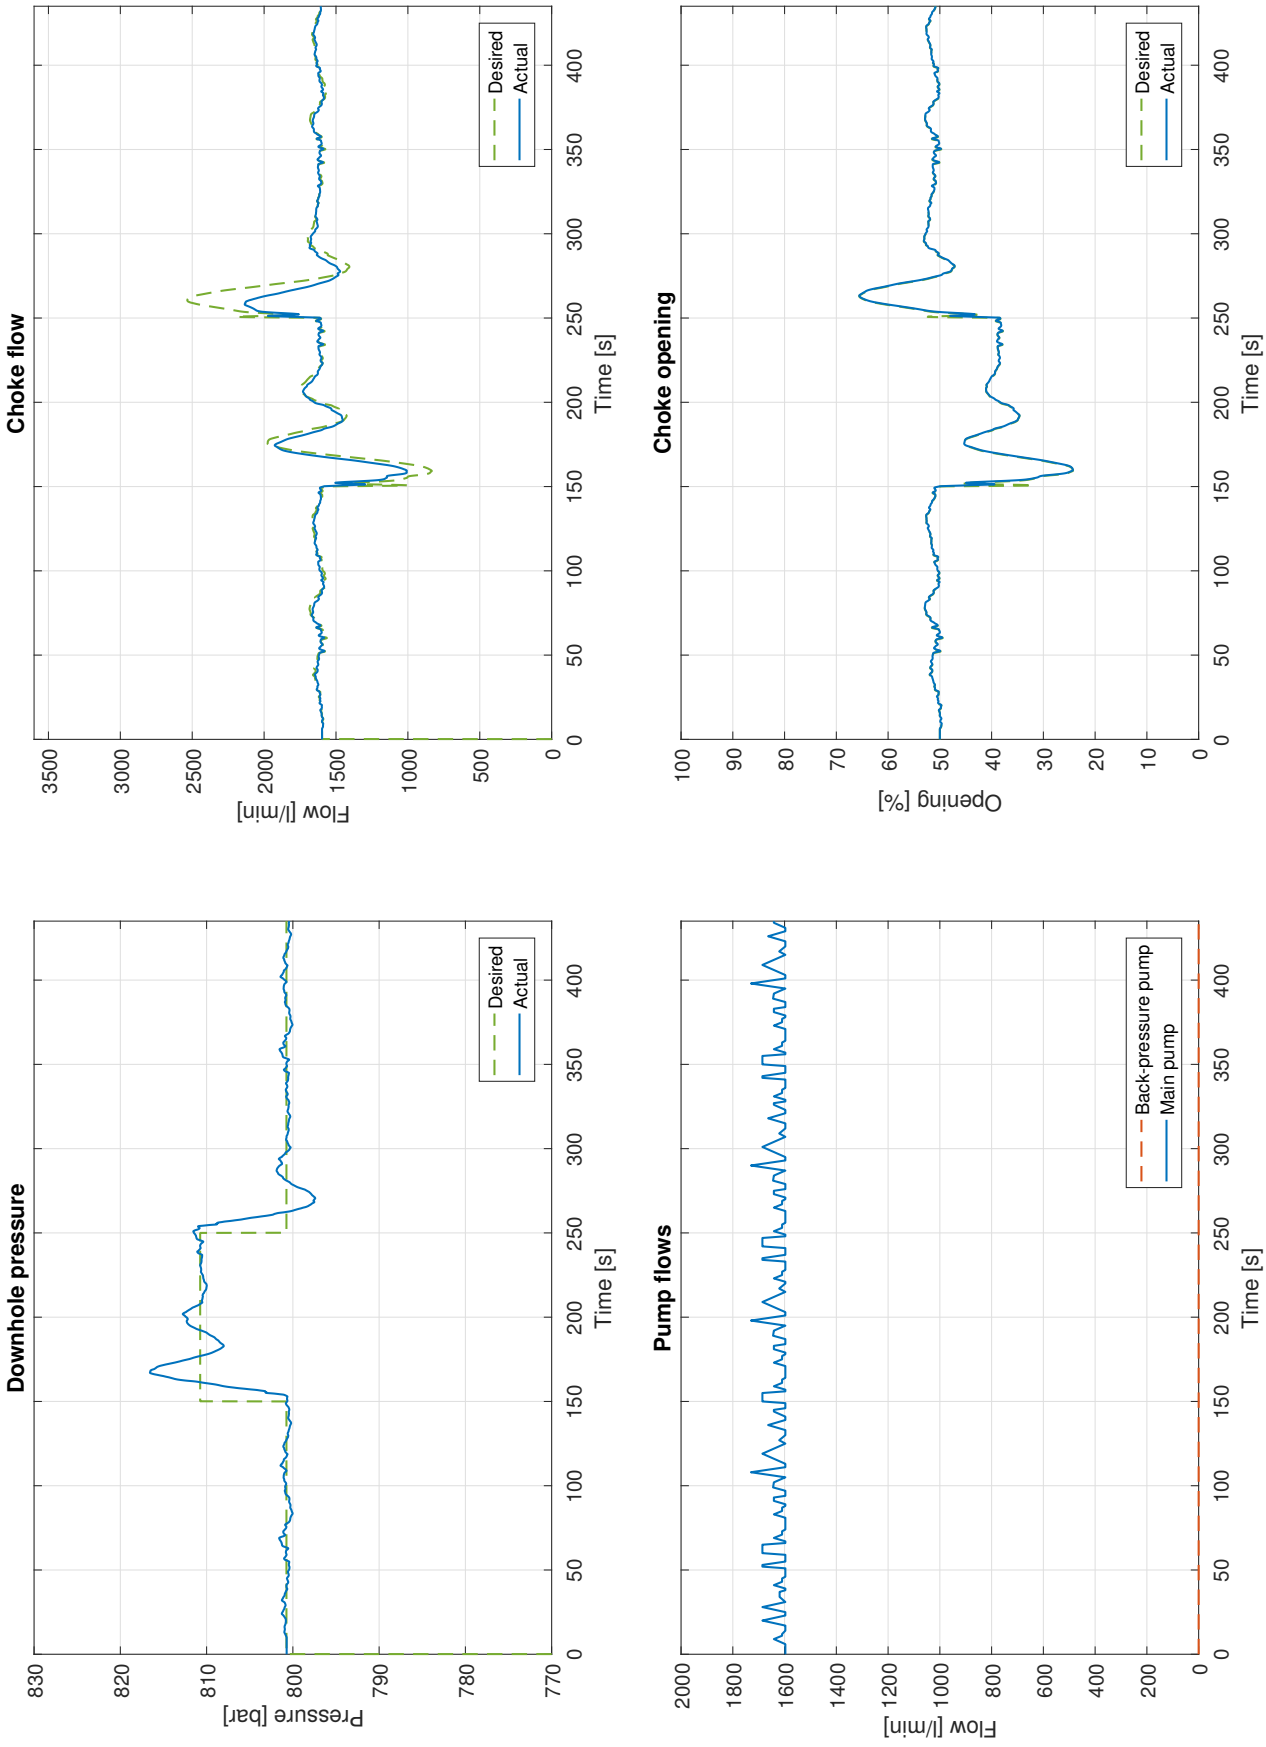

Figure C.4: Setpoint step response with CCS (filtered reference)

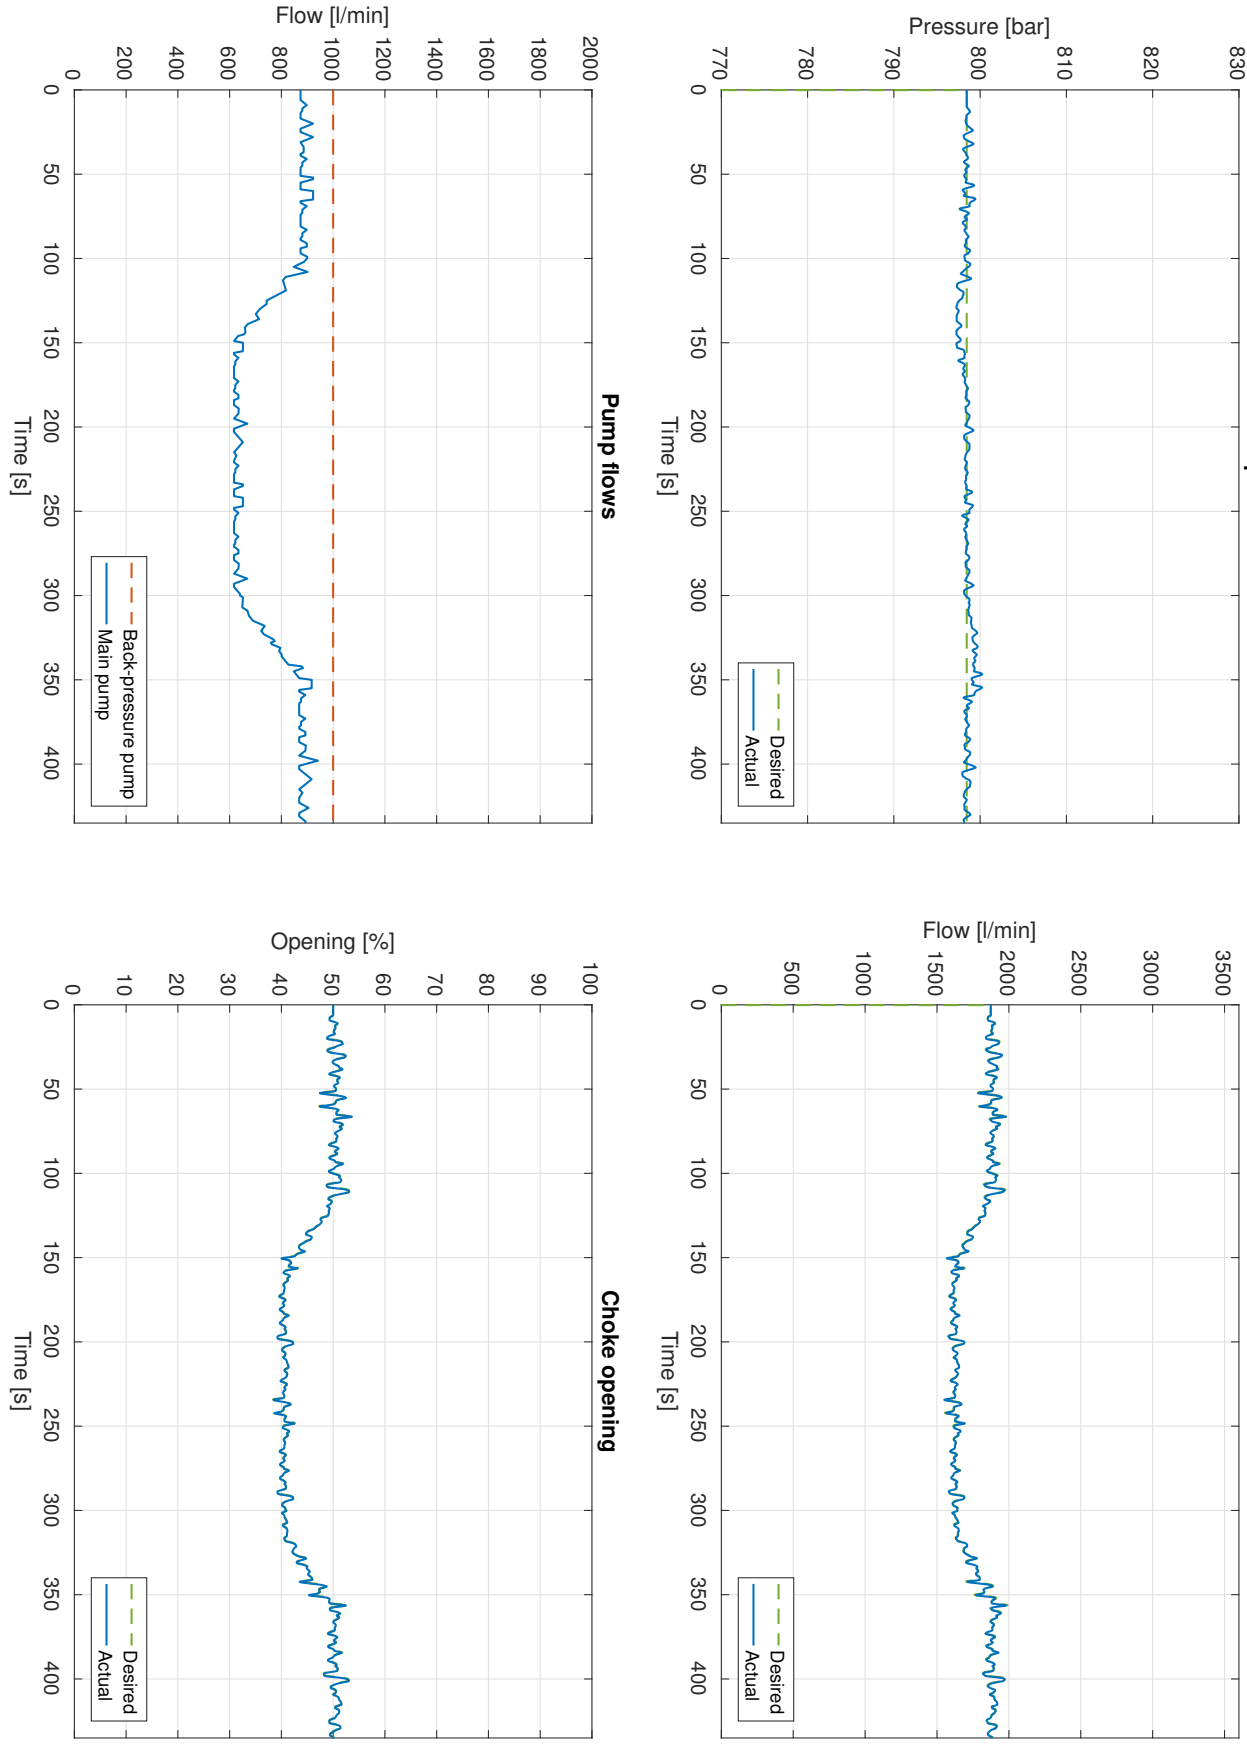

Figure C.5: Pump flow ramp response

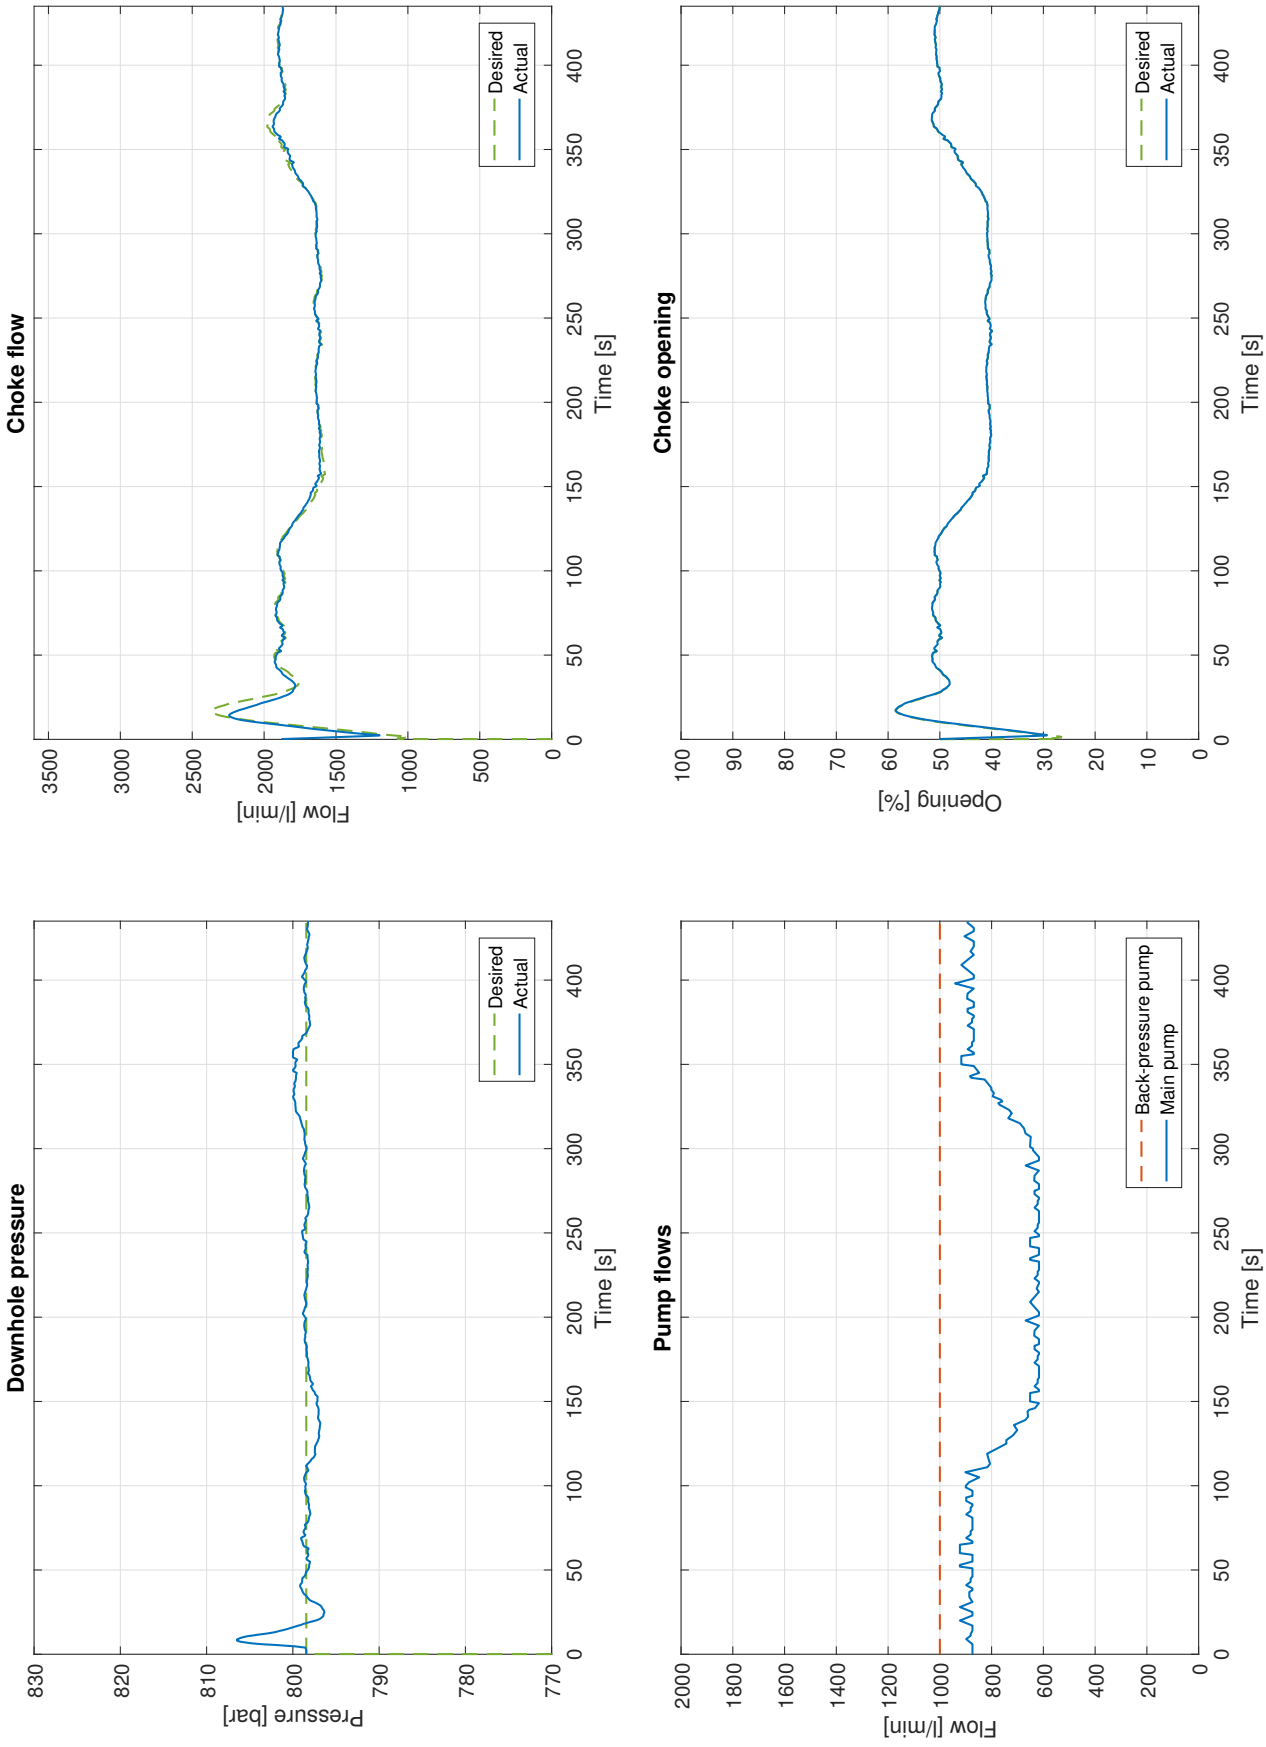

Figure C.6: Pump flow ramp response (filtered reference)

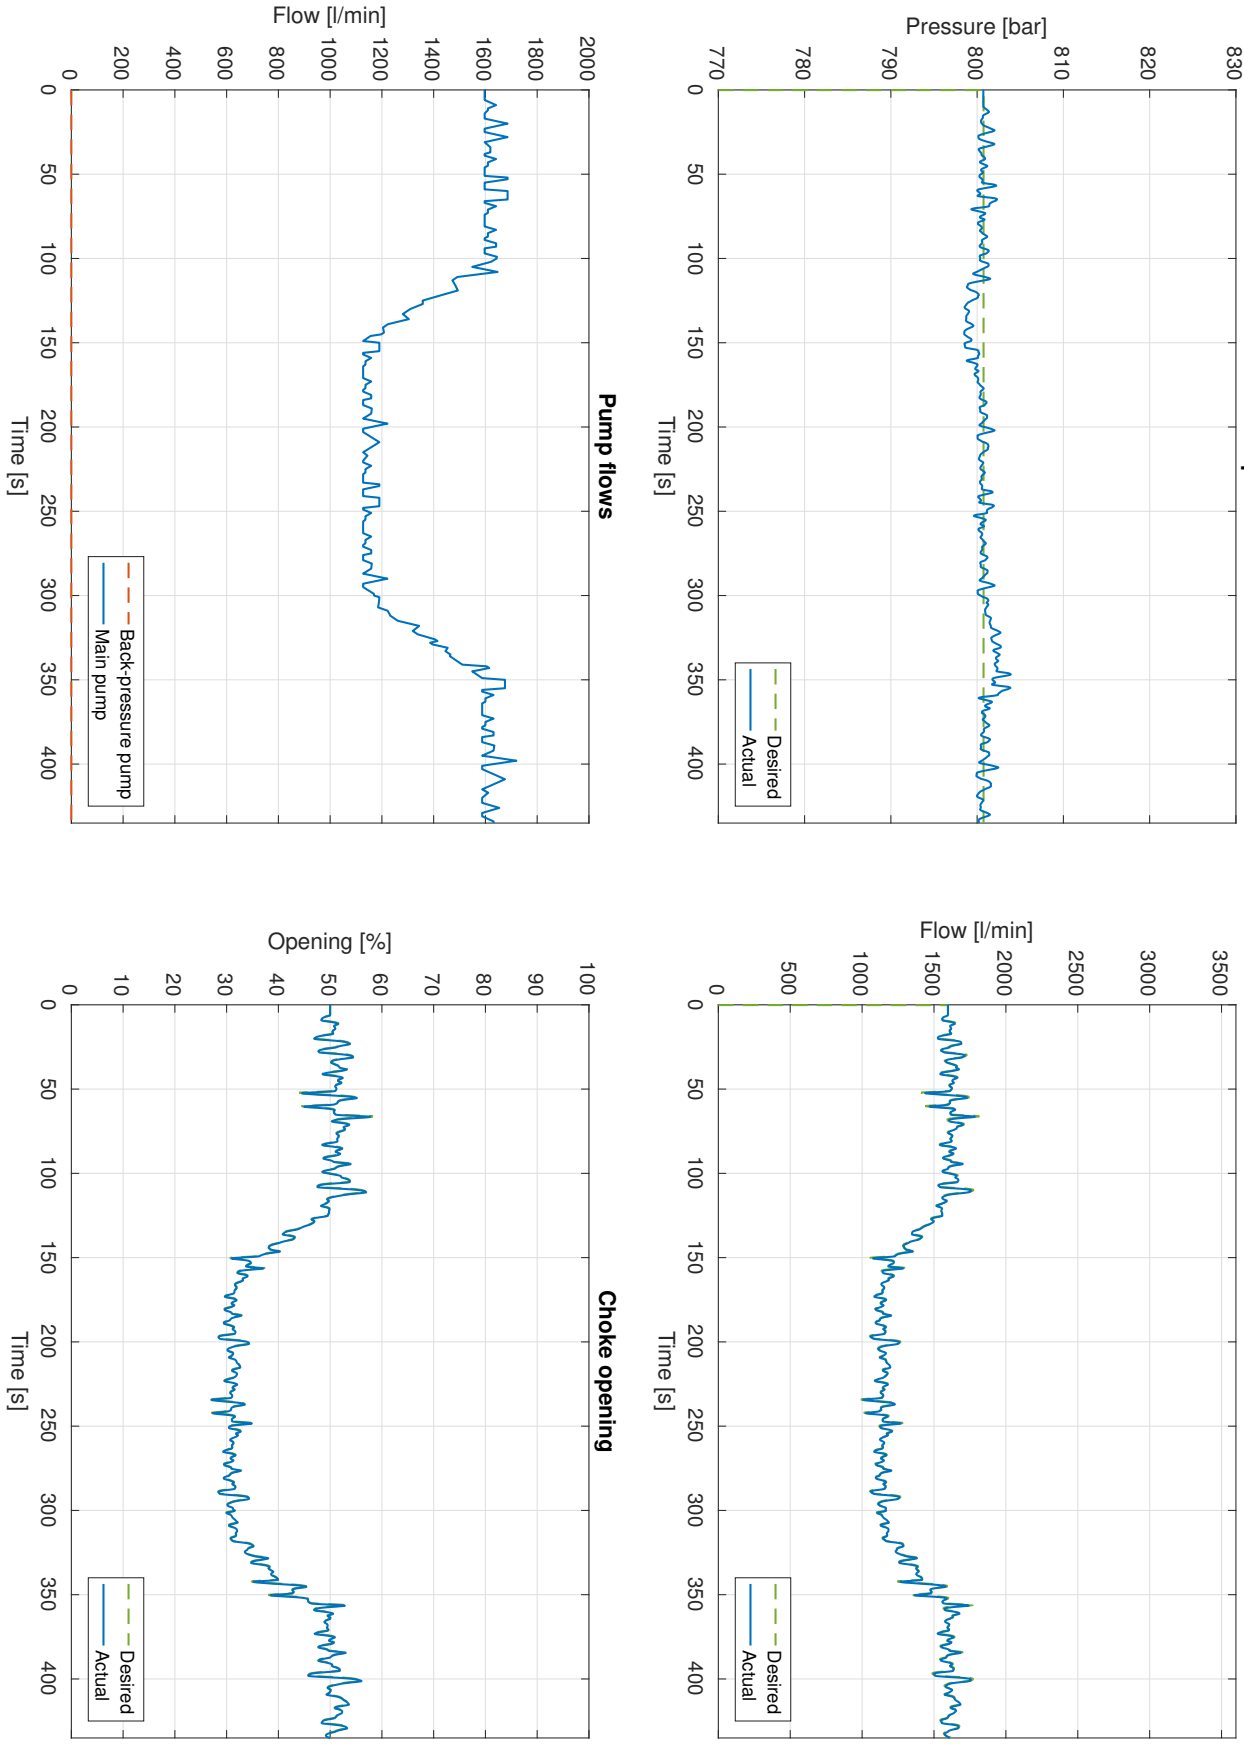

Figure C.7: Pump flow ramp response with CCS

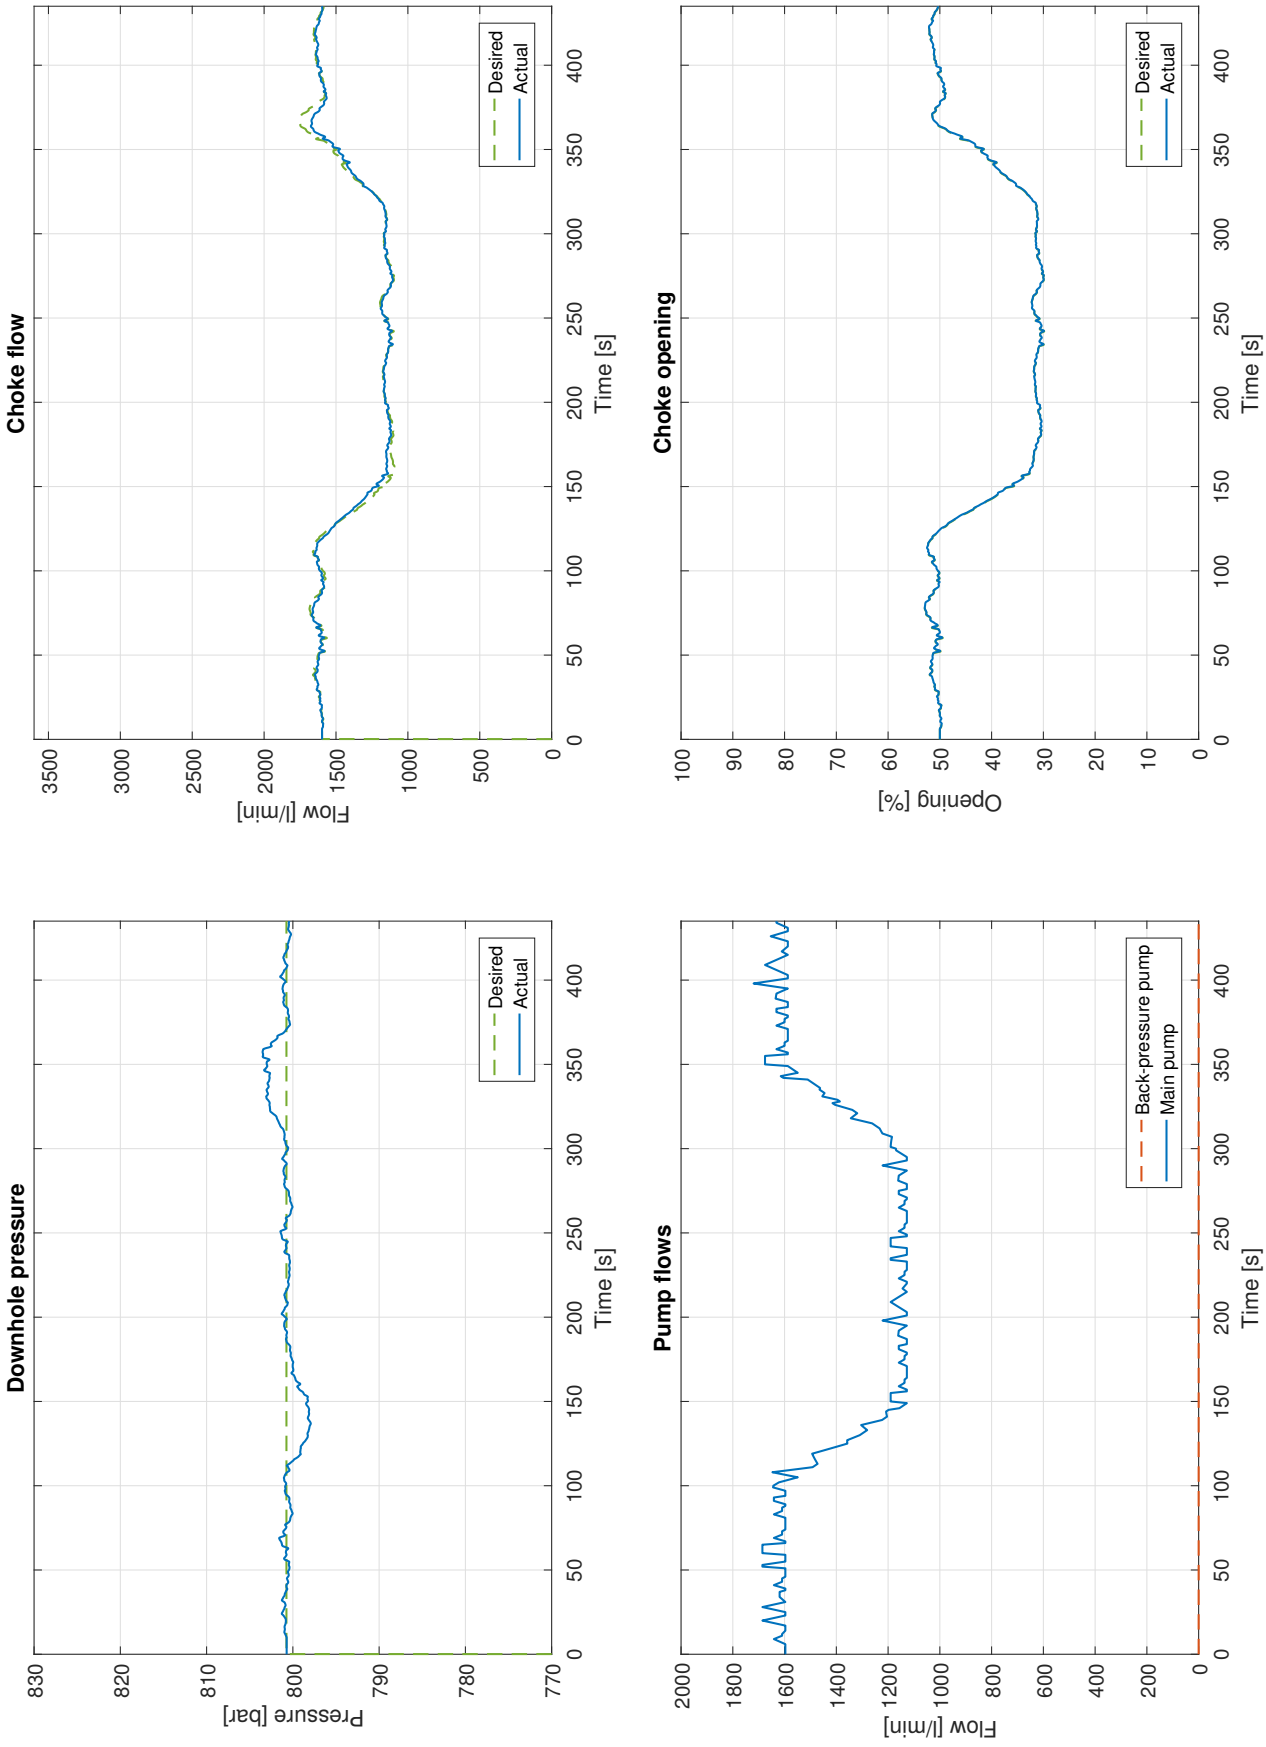

Figure C.8: Pump flow ramp response with CCS (filtered reference)

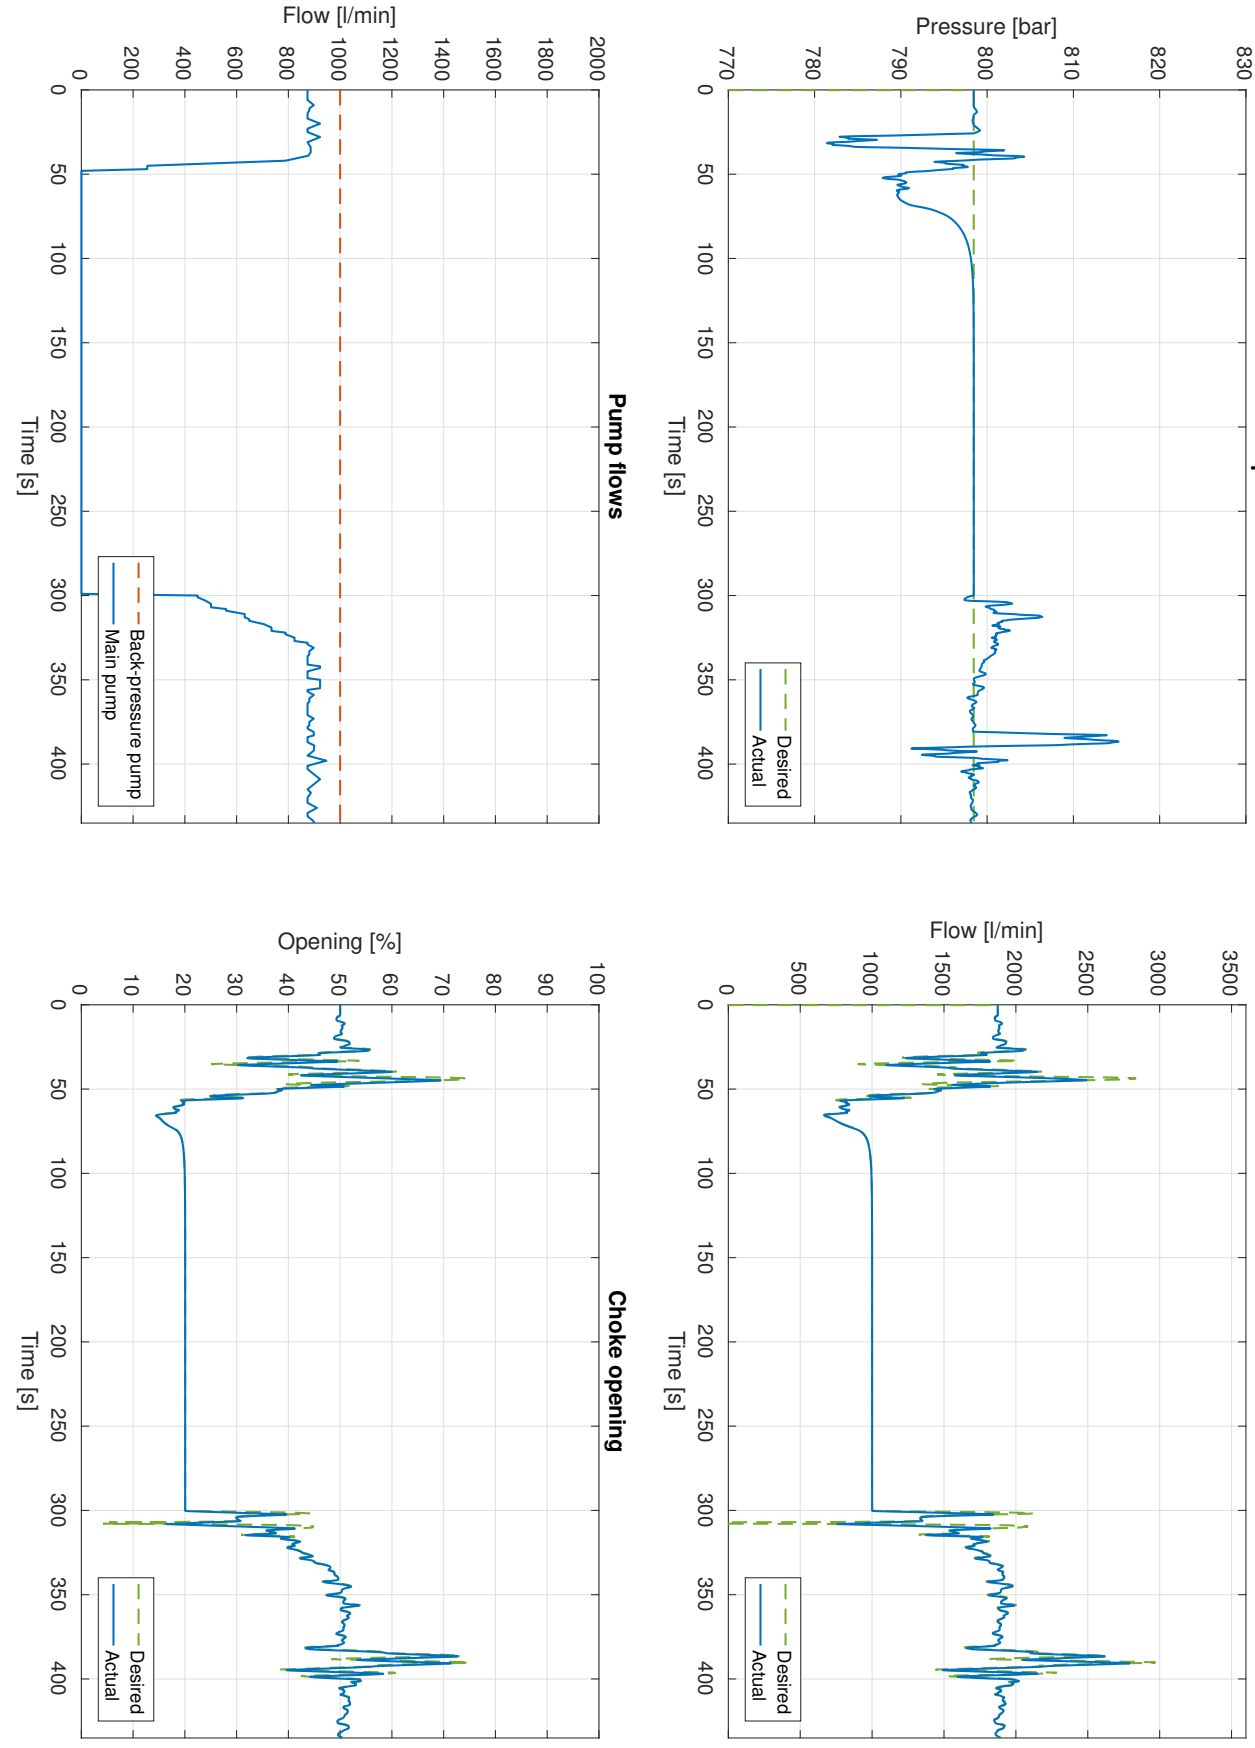

Figure C.9: Connection response in calm conditions

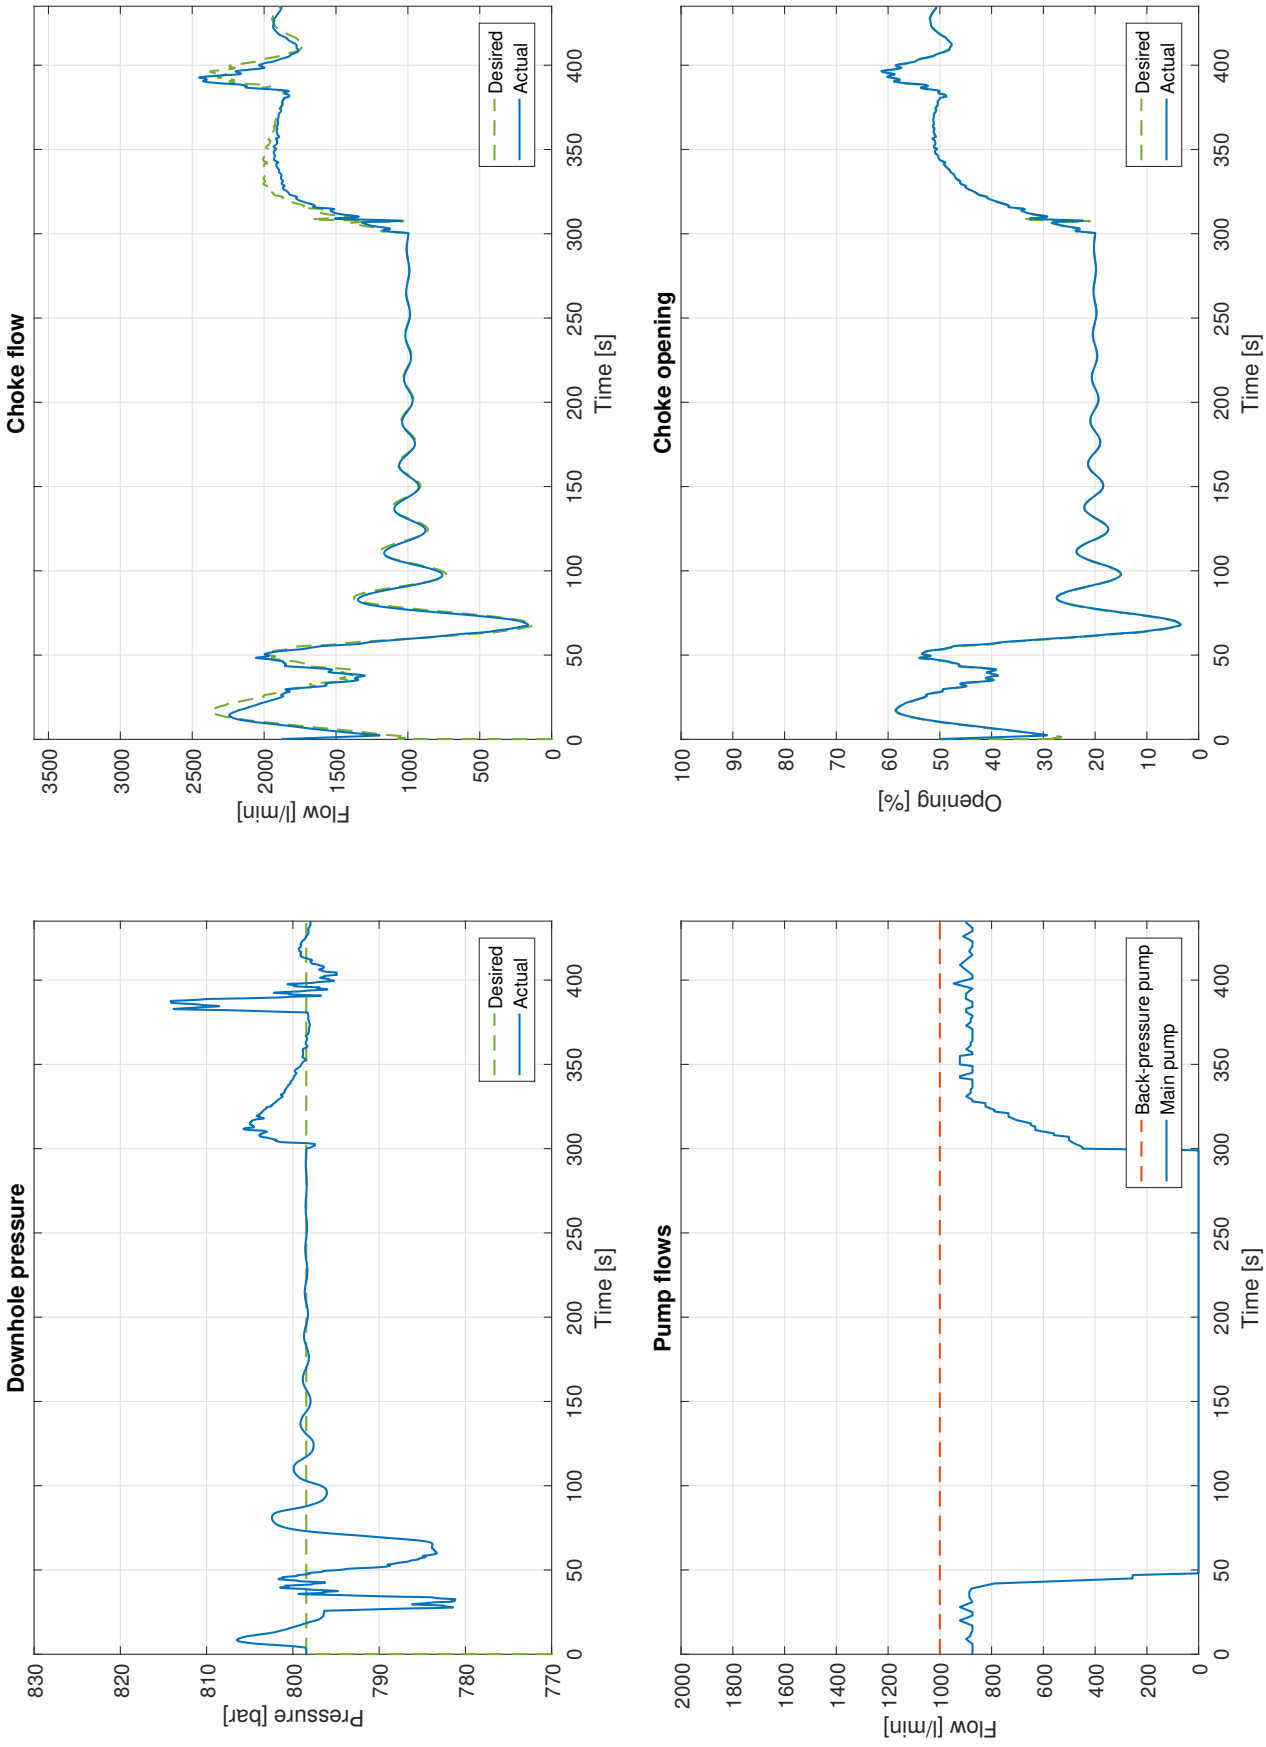

Figure C.10: Connection response in calm conditions (filtered reference)

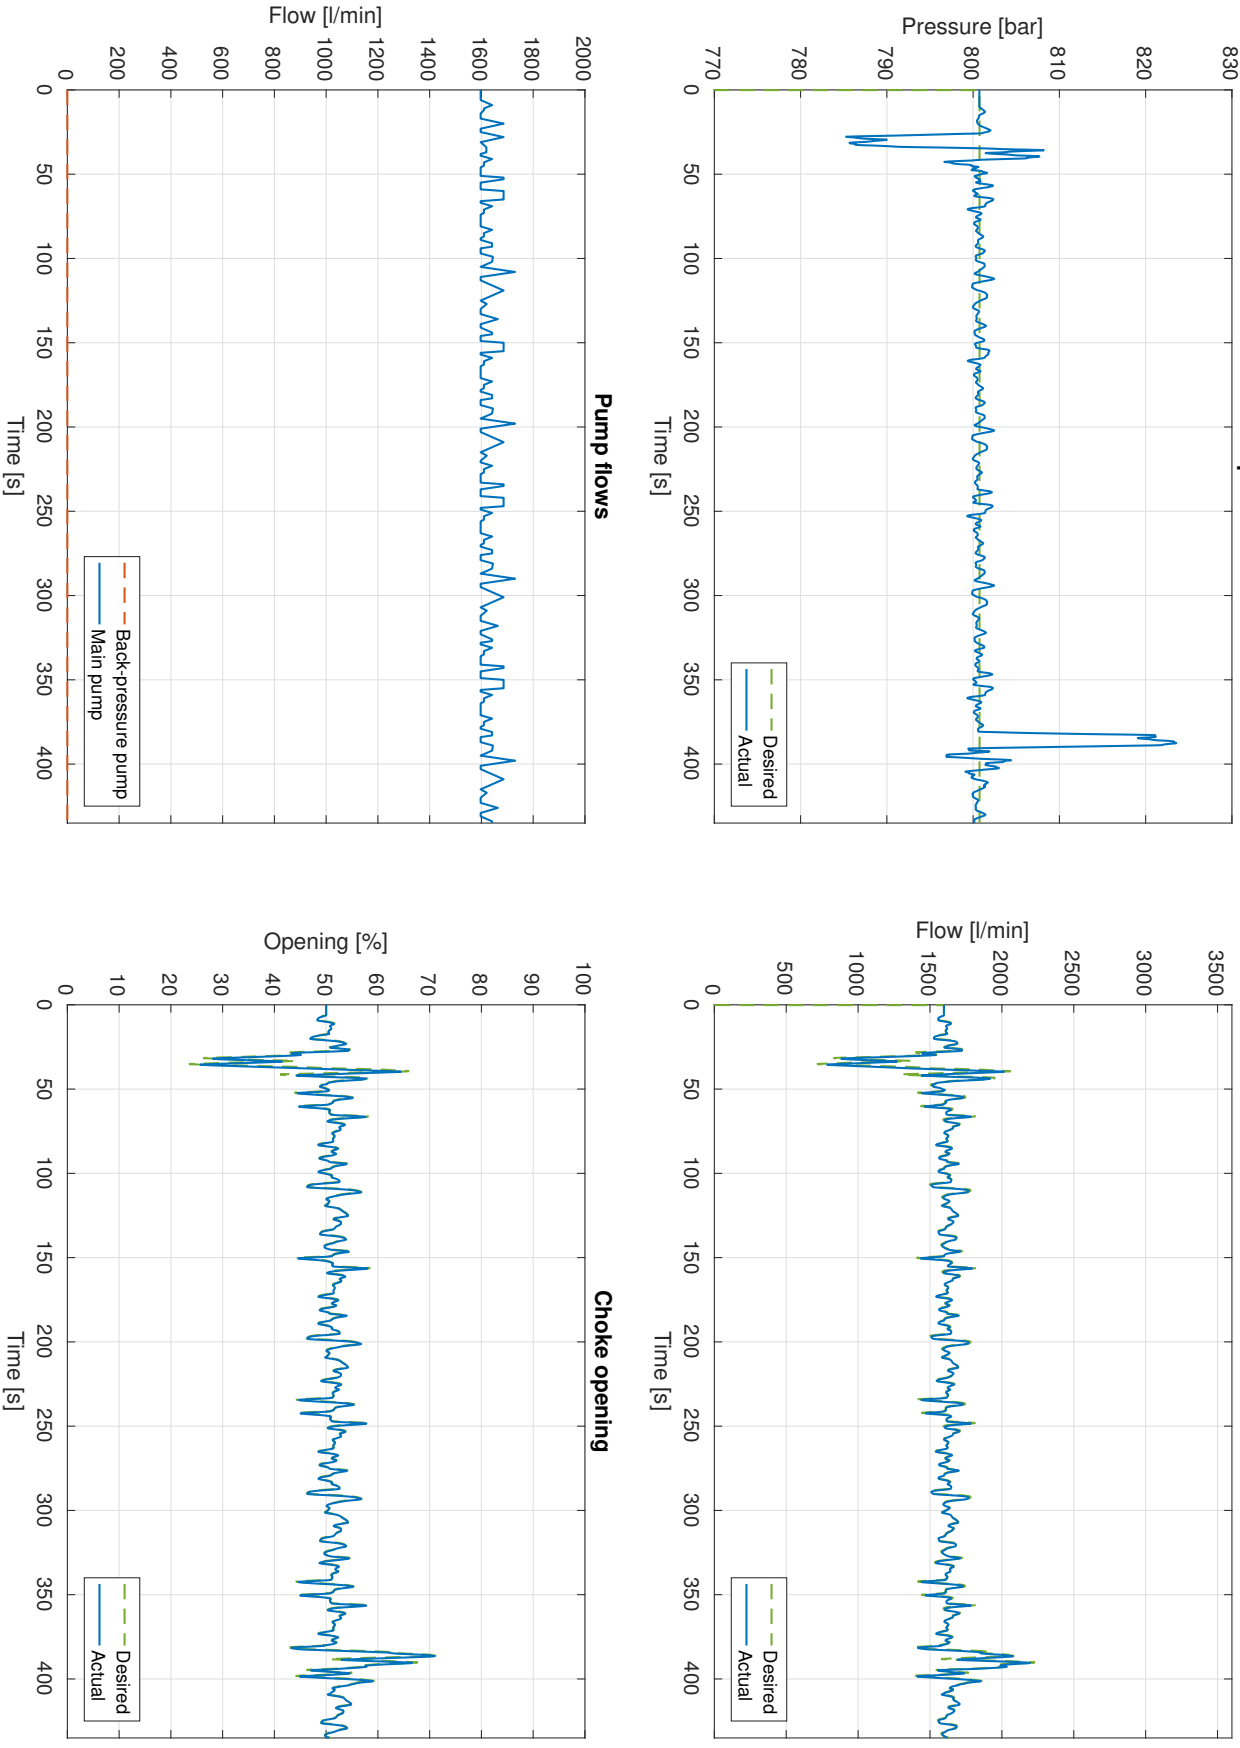

Figure C.11: Connection response in calm conditions with CCS

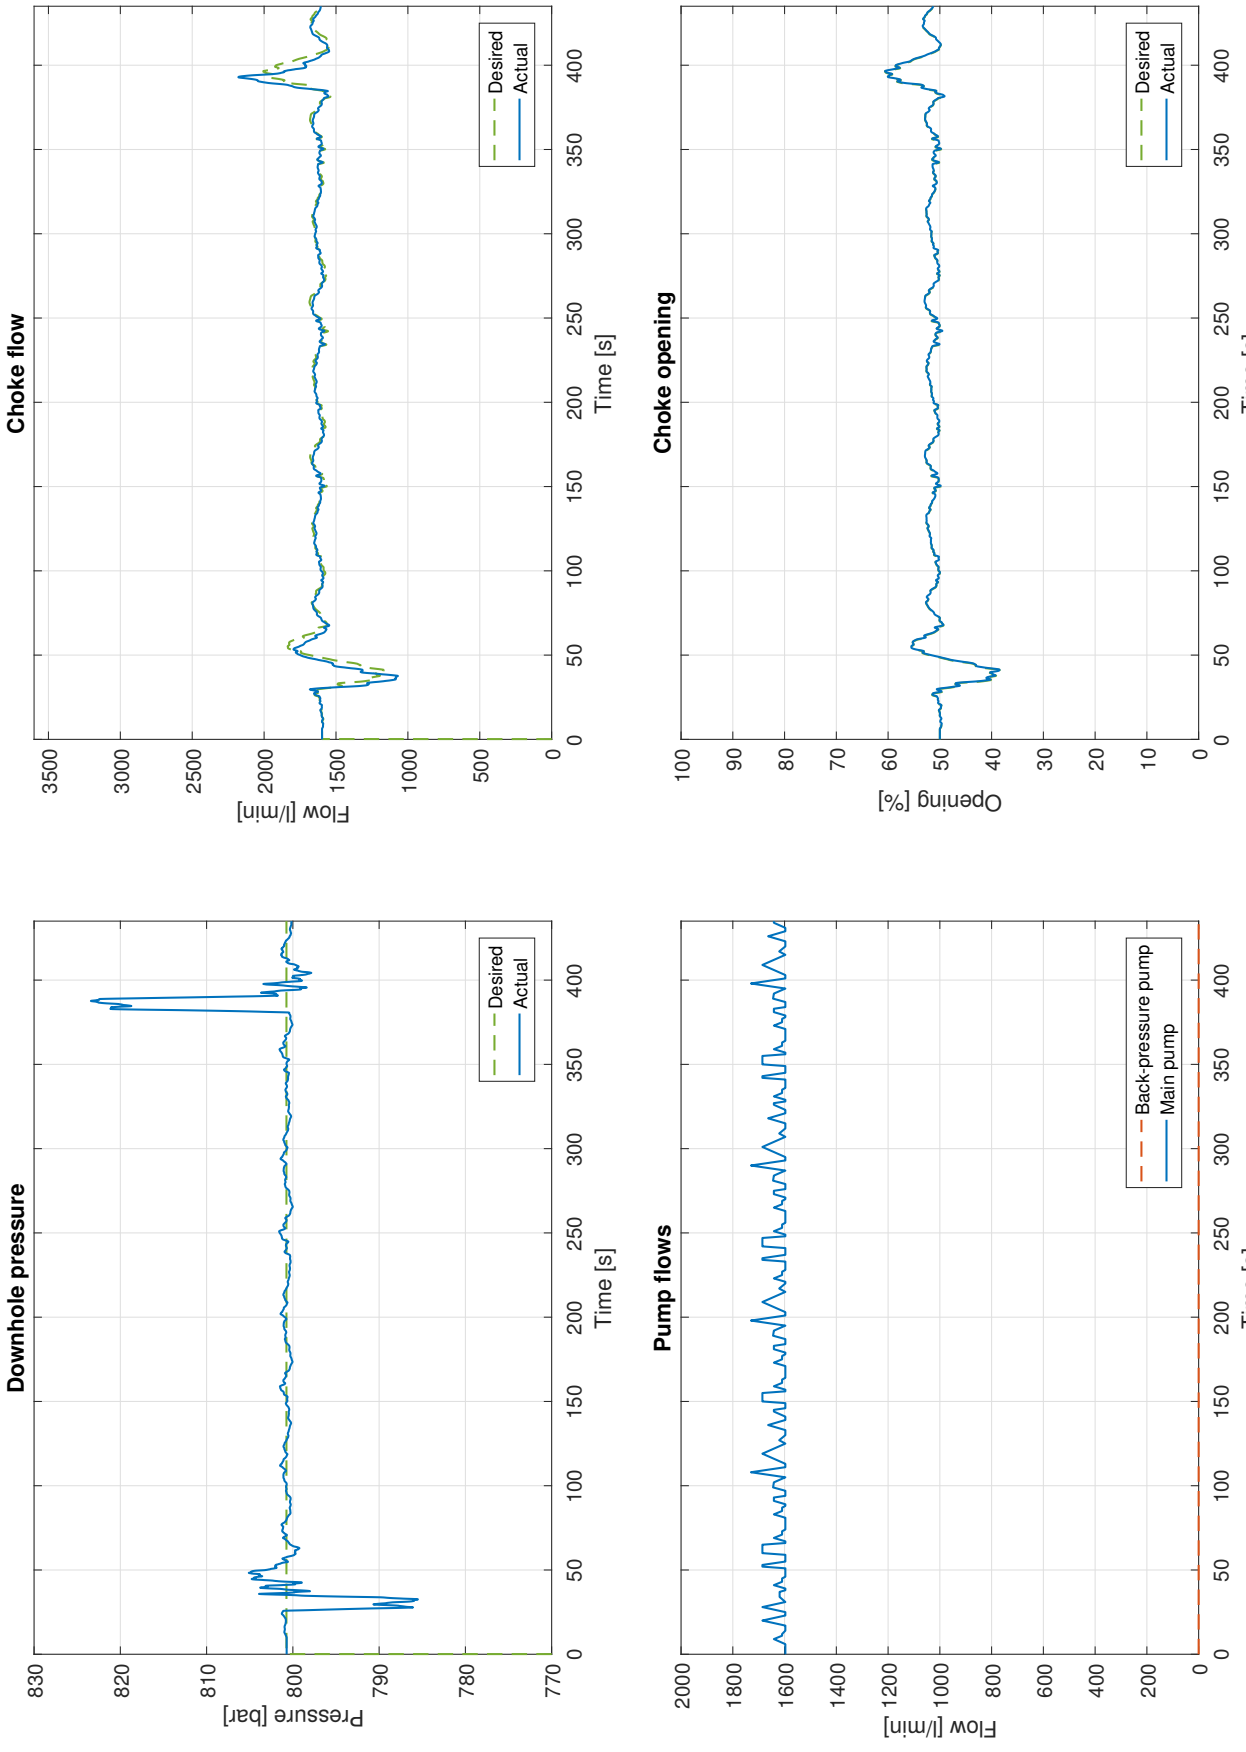

Figure C.12: Connection response in calm conditions with CCS (filtered reference)

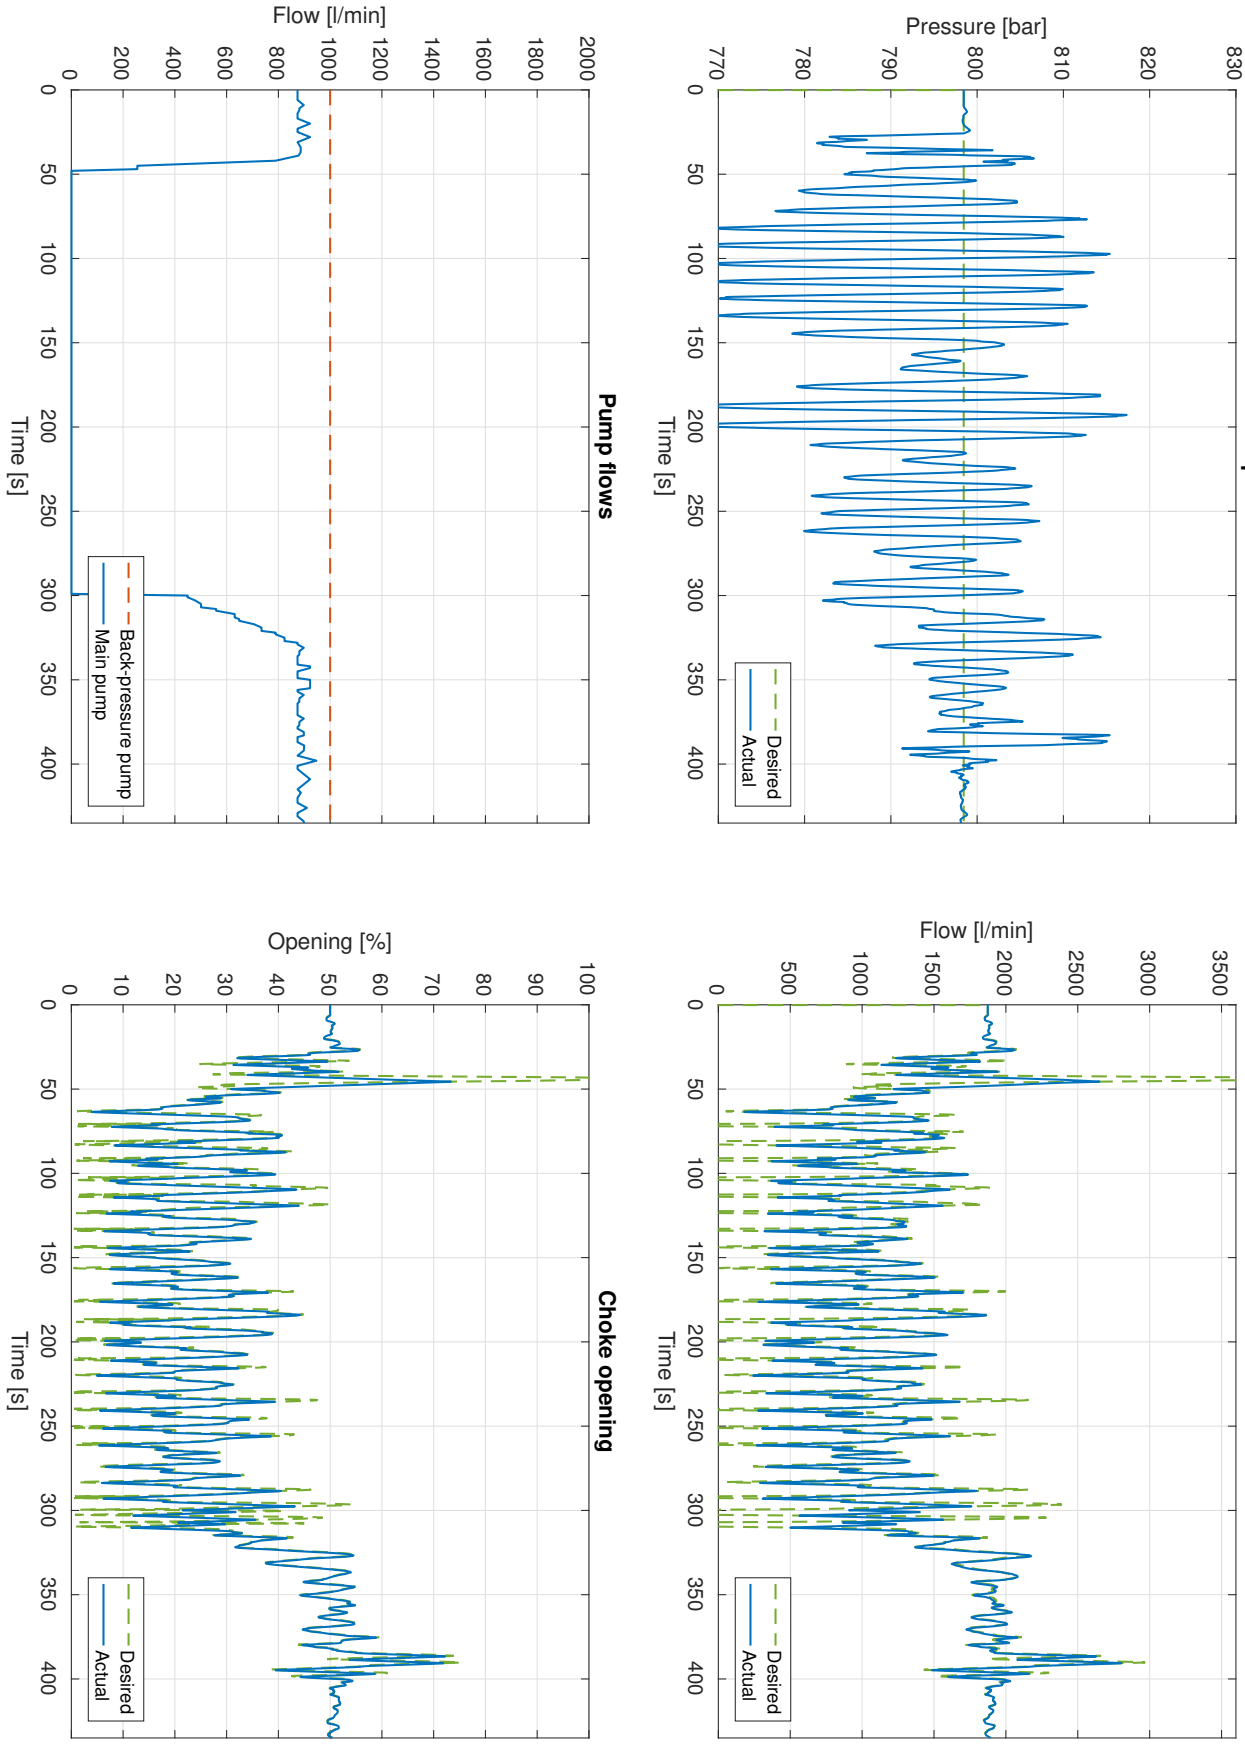

Figure C.13: Connection response in rough conditions

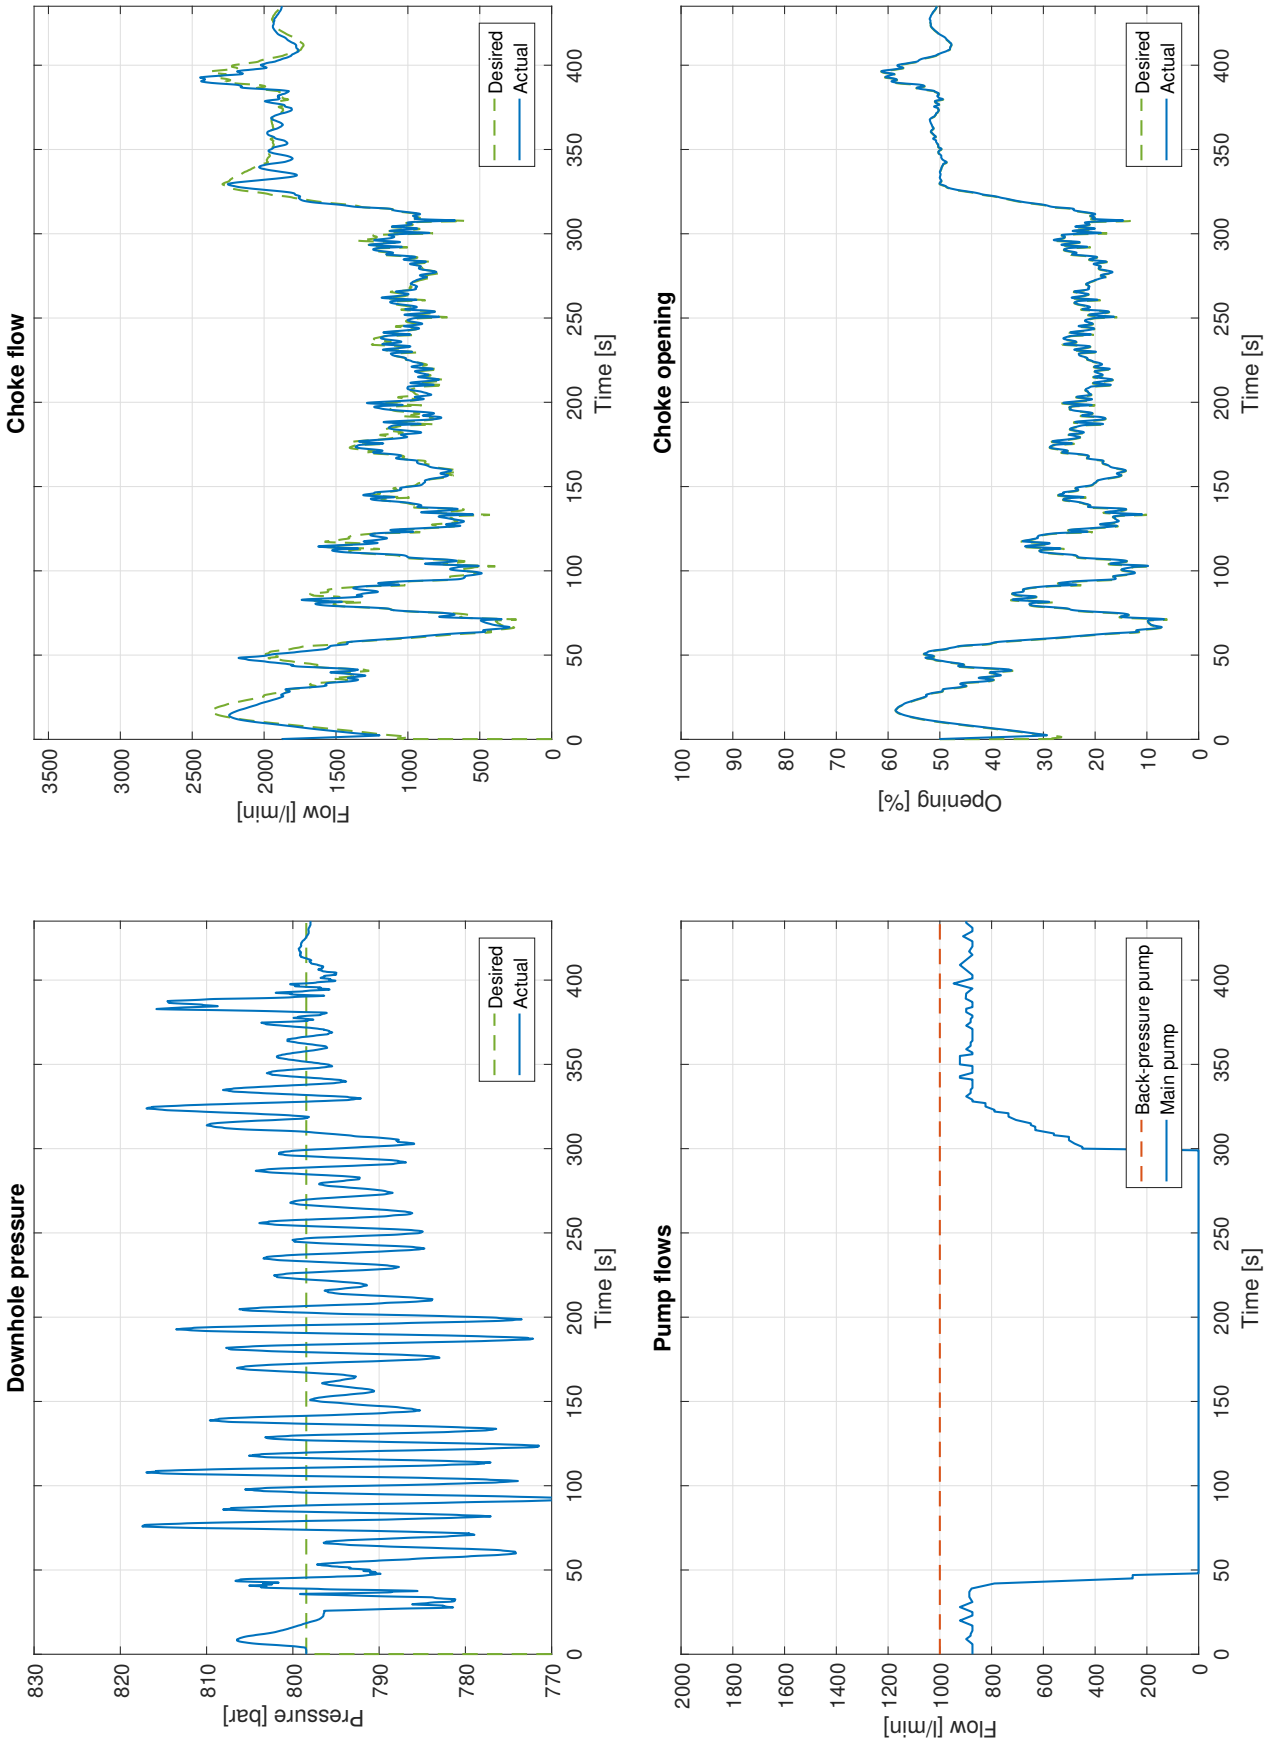

Figure C.14: Connection response in rough conditions (filtered reference)

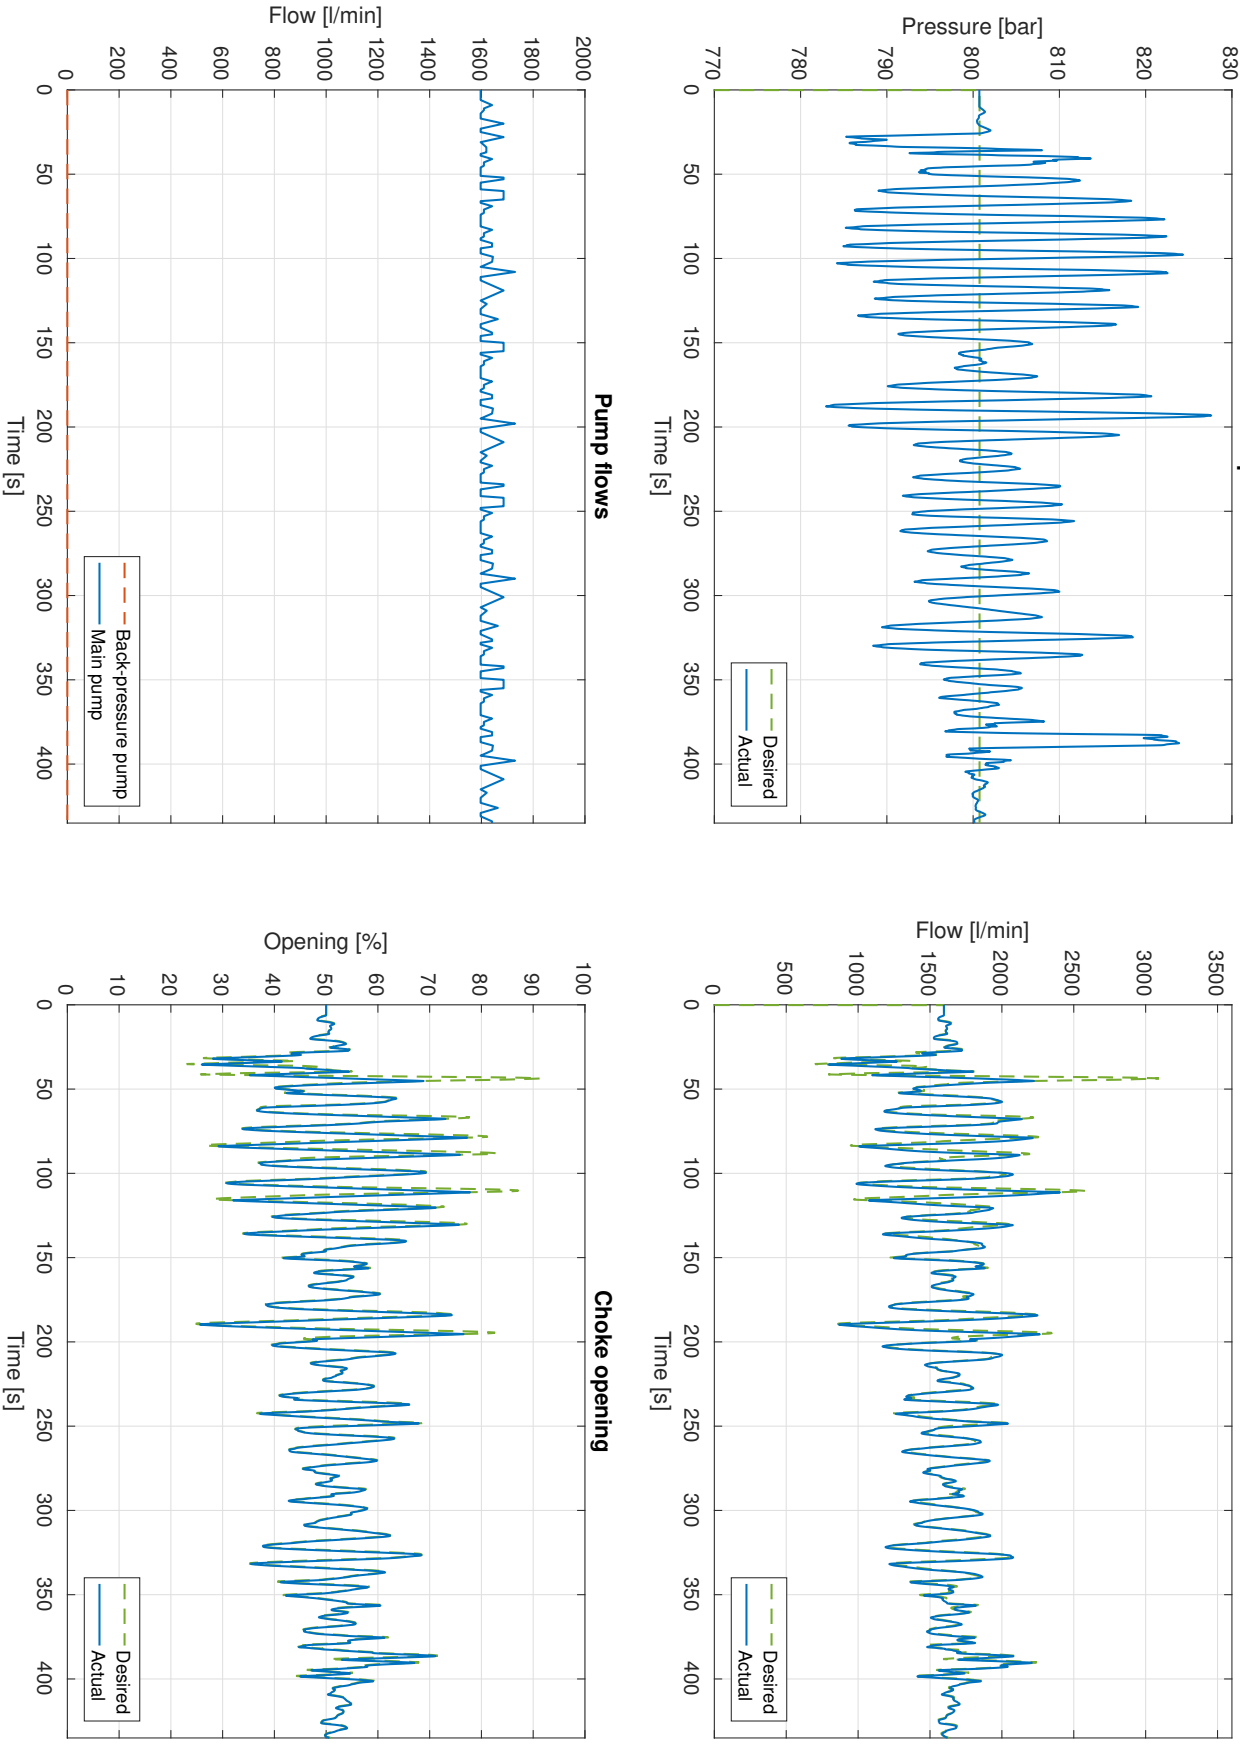

Figure C.15: Connection response in rough conditions with CCS

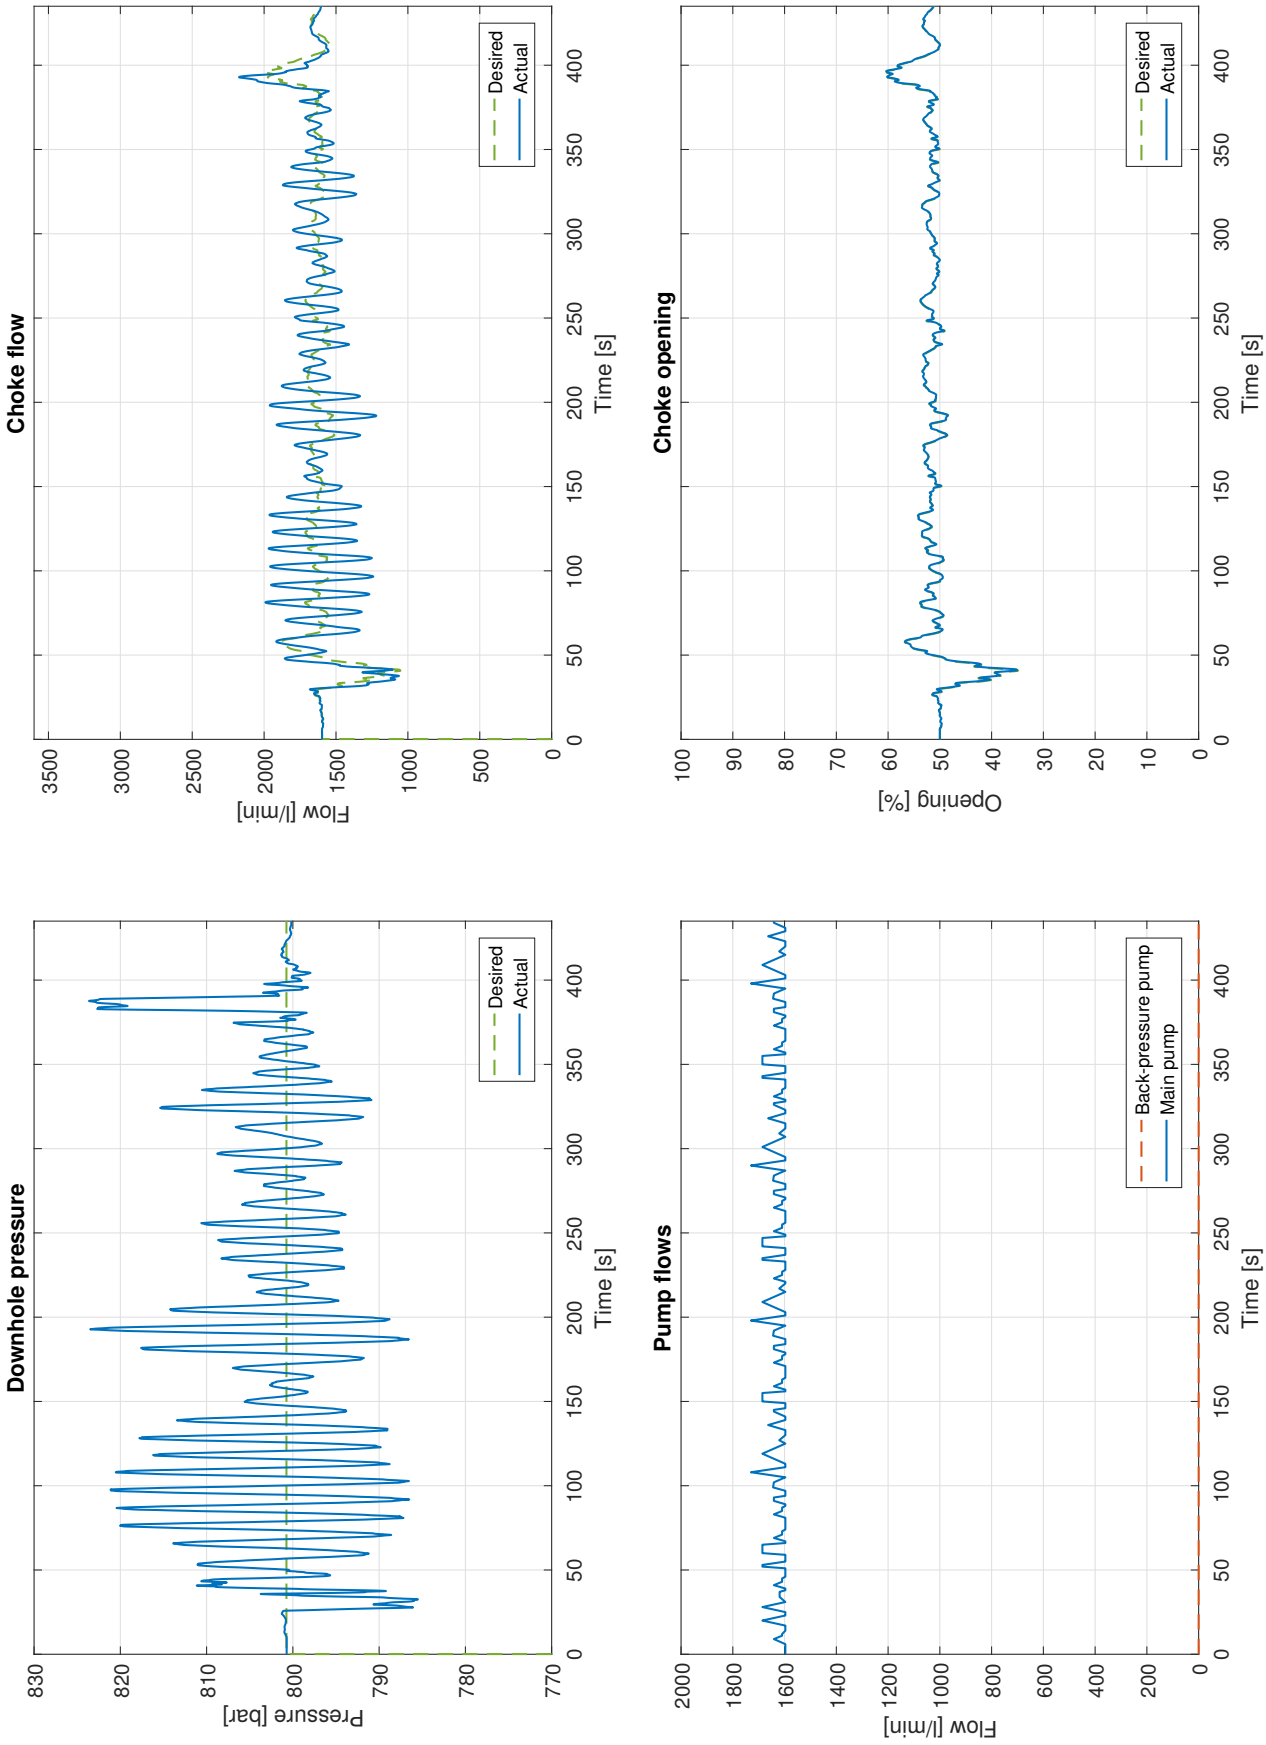

Figure C.16: Connection response in rough conditions with CCS (filtered reference)

References

[1] Unknown. (2009). Drilling the Undrillable, Offshore Technology, [Online]. Available: https://www.offshore-technology.com/features/feature52912 (visited on 06/02/2019).

[2] R. J. Howe, "Evolution of Offshore Drilling and Production Technology", Offshore Technology Conference, Tech. Paper, 1986.

[3] Unknown. (2015). History of offshore drilling units, Society of Petroleum Engineers, [Online]. Available: https://www.offshore-technology.com/features/feature52912 (visited on 06/02/2019).

[4] T. S. Mollan, "An Overview of Managed Pressure Drilling Systems with Implications of Incorporating HeaveLock", Norwegian University of Science and Technology, Project report, 2018.

[5] B. Dow, J. Baker, P. Spriggs, and A. Voshall, "Anatomy of MPD Well Failures: Lessons Learned from a Series of Difficult Wells", Society of Petroleum Engineers/International Association of Drilling Contractors, Tech. Paper, 2013.

[6] D. Elliot, J. Montilva, P. Francis, D. Reitsma, J. Shelton, and V. Roes, "Managed Pressure Drilling Erases the Lines", *Oilfield Review*, vol. 23, no. 1, pp. 14–23, 2011.

[7] T. Strecker and O. M. Aamo, "Limitations of Topside Actuation for Rejecting Heave-induced Pressure Oscillations in Offshore Drilling", Norwegian University of Science and Technology, Tech. Paper, 2018.

[8] A. Albert, O. M. Aamo, J.-M. Godhavn, and A. Pavlov, "Surpressing Pressure Oscillations in Offshore Drilling: Control Design and Experimental Results", Norwegian University of Science and Technology, Tech. Paper, Jul. 2014.

[9] M. Ø. Christensen and E. Gundersen, "Development and Implementation of the Control System for an Autonomous Choke Valve for Downhole Flow and Pressure Control", Master's thesis, Norwegian University of Science and Technology, Trondheim, Norway, 2016.

[10] U. J. F. Aarsnes, O. M. Aamo, E. Hauge, and A. Pavlov, "Limits of Controller Performance in the Heave Disturbance Attentuation Problem", Norwegian University of Science and Technology, Tech. Paper, 2013.

[11] S. Schaut, "Modeling, Estimation and Control for an experimental Lab Facility for Drilling", Master's thesis, University of Stuttgart, 2015.

[12] J. F. H. Totland, "Fast Pressure Control in Managed Pressure Drilling", Master's thesis, Norwegian University of Science and Technology, Trondheim, Norway, 2014.

[13] M. Z. Bardaj, S. Olfati, A. Rafieefar, H. Garmsiri, and E. Zarei, "Continuous Circulation System: A Key to drilling safety increment", Tech. Paper, 2016.

[14] J.-M. Godhavn, A. Pavlov, G.-O. Kaasa, and N. L. Rolland, "Drilling seeking automatic control solutions", Statoil Research Center, Tech. Paper, 2011.

[15] Unknown. (2019). Science of Sound Tutorial, University of Rhode Island/Inner Space Center, [Online]. Available: https://dosits.org/tutorials/science/ (visited on 06/02/2019).

[16] ——, (2019). Sound is a Pressure Wave, The Physics Classroom, [Online]. Available: https://www.physicsclassroom.com/class/sound/Lesson-1/Sound-is-a-Pressure-Wave (visited on 06/02/2019).

[17] H. F. George and A. Barber. (2007). Bulk Modulus: What is it? When is it important?, The Lubrizol Corp., [Online]. Available: https://www.hydraulicspneumatics.com/200/TechZone/HydraulicFluids/Article/False/70094/TechZone-HydraulicFluids (visited on 06/02/2019).

[18] K. Bjørvik, *Dynamiske systemer*.

[19] T. I. Fossen, *Handbook of Marine Craft Hydrodynamics and Motion Control*. John Wiley & Sons, Ltd, 2011.

[20] M. J. Grimble and M. A. Johnson, *Optimal Control and Stochastic Estimation: Theory and Applications*. John Wiley & Sons, Ltd, 1989.

[21] K. Torsethaugen, "Model for a Doubly Peaked Wave Spectrum", SINTEF, Tech. Paper, 1996.

[22] O. Egeland and T. Gravdahl, *Modeling and Simulation for Automatic Control*. Marine Cybernetics AS, 2003.

[23] M. A. B. Yusof and S. Kabir, "An Overview of Sonar and Electromagnetic Waves for Underwater Communication", IETE, Tech. Paper, 2012.

[24] M. J. Jellison, D. R. Hall, D. C. Howard, J. H. Tracy Hall, R. C. Long, R. B. Chandler, and D. S. Pixton, "Telemetry Drill Pipe: Enabling Technology for the Downhole Internet", Society of Petroleum Engineers, Tech. Paper, 2003.

[25] H. YoonSong and H.-g. Kim, "A Real-time Filtering Method of Positioning Data with Moving Window Mechanism", Hongik University, Tech. Paper, 2012.

[26] V. Czitrom and P. D. Spagon, *Statistical Case Studies for Industrial Process Improvement*. 1997.

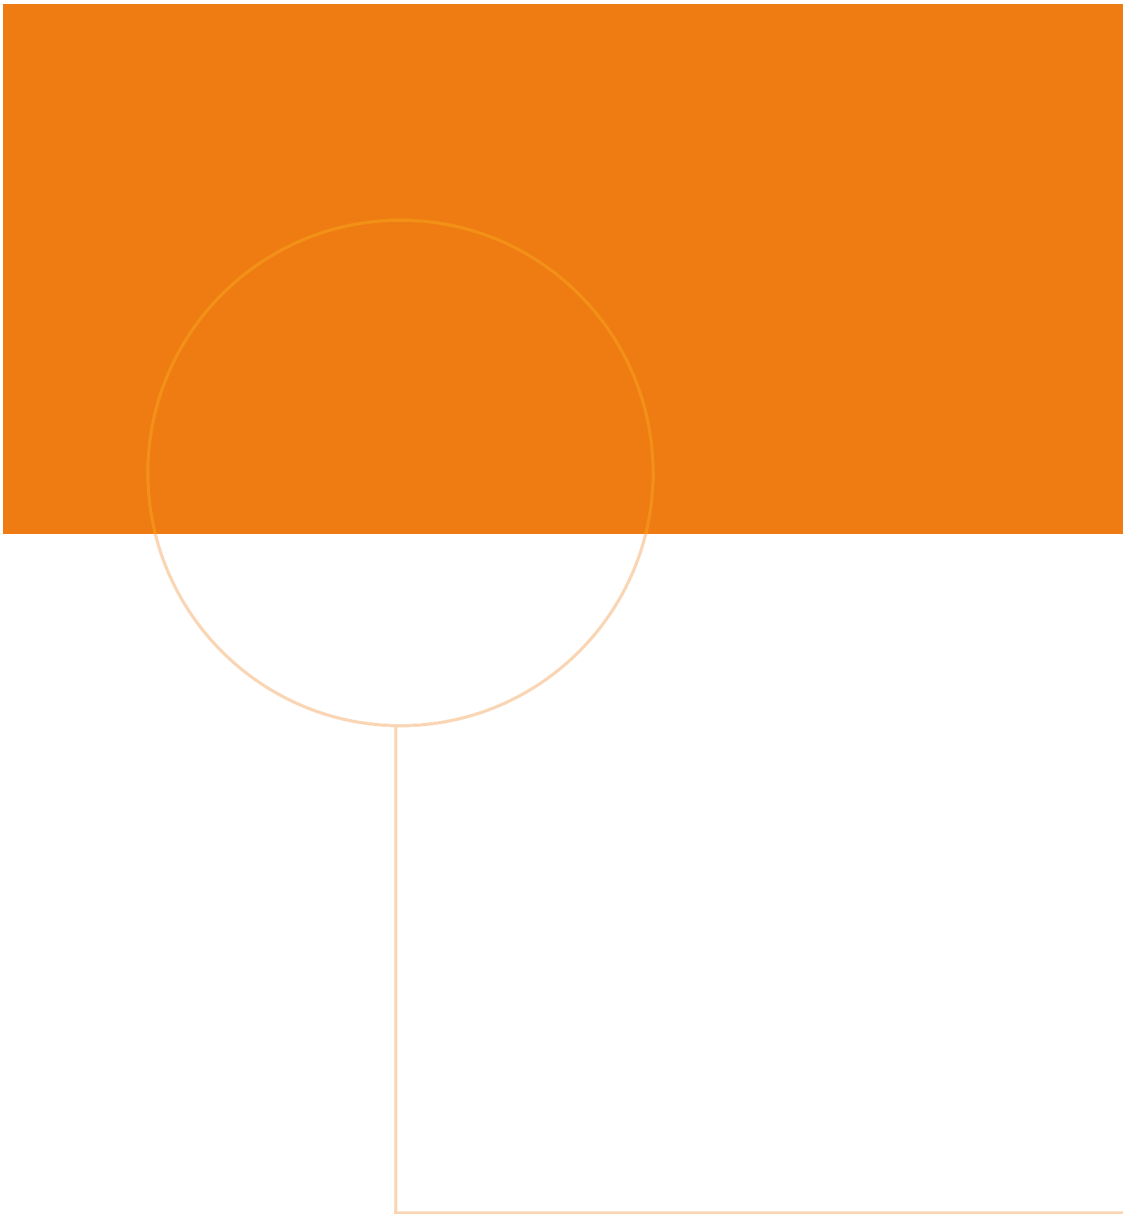

NTNU
Norwegian University of
Science and Technology