

https://dspace.org/

The Control of Flotation Cells at the Point of Peak Air Recovery

Ramodike, Katlego Samantha

2025-10
University of the Witwatersrand, Johannesburg
https://hdl.handle.net/10539/49864
Ramodike, Katlego Samantha. (2025). The Control of Flotation Cells at the Point of Peak Air Recovery. [Master's dissertation, University of the Witwatersrand, Johannesburg]. WIReDSpace. https://hdl.handle.net/10539/49864

Downloaded from a DSpace repository. https://dspace.org/



UNIVERSITY OF THE
WITWATERSRAND,
JOHANNESBURG

The Control of Flotation Cells at the Point of Peak Air Recovery

Master of Engineering (Chemical Engineering)

Prepared by:

Katlego S. Ramodike

600873

Submitted to:

School of Chemical and Metallurgical Engineering, Faculty of Engineering and the Built Environment, University of the Witwatersrand,

Johannesburg, South Africa.

Supervisors: Professor Kevin Brooks, Dr Antony Higginson

October 2025

ABSTRACT

Flotation efficiency is closely linked to air recovery, making its optimisation critical for improved stability and metallurgical performance. This study presents a proportional (P) control strategy designed to maintain a flotation cell at Peak Air Recovery (PAR). The control approach was developed, implemented, and evaluated through numerous simulations, with performance quantified over a range of operating conditions

Under noise free conditions, the controller achieved PAR within approximately 25 seconds, demonstrating fast and stable convergence. In moderate noise cases, the controller maintained air recovery with minimal fluctuations, remaining close to the target PAR value adapting effectively to process fluctuations. The study also investigated disturbance rejection, with the controller converging to PAR within 90 seconds following abrupt changes in superficial gas velocity. However, disturbances that shifted the process optimum, such as froth height increases, could not be fully counteracted by manipulating the gas velocity alone.

The controller's performance was compared with Extremum Seeking Control (ESC), while adaptive, converged more slowly under similar conditions. The study also touched on the use of Model Predictive Control (MPC), noted for its advanced approach in dynamically adjusting operational parameters to optimise flotation efficiency. Although the specific performance details like the speed of achieving PAR with MPC were not fully explored, its inclusion highlights the scope of current control methodologies.

The findings demonstrate that a well tuned P controller offers a simple, fast and reliable means of maintaining operation at PAR, providing an alternative to more complex optimisation based strategies.

Keywords: Flotation, peak air recovery (PAR), process control, Proportional control, Extremum Seeking Control (ESC), Model Predictive Control (MPC), flotation optimisation.

ACKNOWLEDGEMENTS

I would like to recognise the contributions of all those who have played a role in making this work possible.

I am deeply grateful to my parents for their unwavering support, encouragement, and sacrifices throughout my academic journey. Their belief in my abilities and constant motivation have been instrumental in my success, and I am forever thankful for their guidance and love.

I extend my sincere appreciation to my supervisors, Professor Kevin Brooks and Dr. Antony Higginson, for their invaluable mentorship, insightful guidance, and continuous support throughout this research. Their expertise, patience, and constructive feedback have been crucial in shaping this dissertation. Through their guidance, I have gained significant analytical and programming skills, as well as a deeper understanding of the subject matter. The knowledge and experience I have acquired during this project has been invaluable, and I am grateful for the opportunity to learn and grow under their supervision.

I would also like to express my heartfelt gratitude to my partner, Gerald Lehana, for his support and encouragement throughout this journey. His understanding, reassurance, and ability to keep me grounded during challenging times have been invaluable. His presence has provided me with the strength and motivation to persevere, and I am truly thankful for his unwavering belief in me.

DECLARATION

I, Katlego S Ramodike, declare that this dissertation is my own, unaided work. It is being submitted for the Master of Science (MSc) degree at the University of the Witwatersrand, Johannesburg. It has not been submitted before for any degree or examination at any other University.



Katlego S. Ramodike

21st day of October 2025 in Krugersdorp

CONTENTS

ABSTRACT II

ACKNOWLEDGEMENTS III

CONTENTS V

LIST OF FIGURES VIII

LIST OF TABLES X

NOMENCLATURE XI

1 INTRODUCTION 1

1.1 PROBLEM STATEMENT 3
1.2 RESEARCH AIM AND OBJECTIVES 4
1.2.1 RESEARCH AIM 4
1.2.2 RESEARCH OBJECTIVES 4

2 LITERATURE REVIEW 6

2.1 THE FUNDAMENTALS OF FROTH FLOTATION 6
2.1.1 OPERATIONAL PARAMETERS 7
2.1.1.1 Air Profiles 7
2.1.1.2 pH Control 8
2.1.1.3 Reagent Type and Dosage 8
2.1.1.4 Agitation Speed 8
2.1.1.5 Particle Size Distribution 9
2.1.2 SAMPLING INTERVALS IN INDUSTRIAL FLOTATION 13
2.2 FROTH STABILITY 14
2.3 AIR RECOVERY 16
2.3.1 FACTORS INFLUENCING AIR RECOVERY 17
2.3.1.1 Air Profiles and Air Distribution 17
2.3.1.2 Froth Height 17
2.4 PEAK AIR RECOVERY 19
2.5 CONTROL THEORY IN FLOTATION PROCESSES 21
2.5.1 REGRESSION METHODS IN CONTROL APPLICATIONS 22
2.5.1.1 Regression Techniques in Control 22
Linear and Multivariate Regression 23
Ridge regression 23
Moore-Penrose Inverse 24
2.5.2 CONTROL STRATEGIES. 24
2.5.3 PROPORTIONAL INTEGRAL DERIVATIVE CONTROL 25

2.5.4 MODEL PREDICTIVE CONTROL 26
2.5.5 EXTREMUM SEEKING CONTROL 28
2.6 REVIEW OF PAST AND CURRENT RESEARCH 31

3 MODELLING METHOD AND SYSTEM DEVELOPMENT 35

3.1 MODELLING METHOD 35
3.1.1 DATA EXTRACTION 37
3.1.2 MODEL VALIDATION 40
3.1.2.1 Dynamic Model 41
3.1.2.2 Steady State Model 41
3.2 PARAMETER ESTIMATION AND REGRESSION ANALYSIS 42
3.2.1 REGRESSION METHOD 42
3.3 SYSTEM DEVELOPMENT 45
3.3.1 REGRESSION APPLICATION IN SIMULINK 47
3.3.2 SETPOINT CALCULATION METHOD 48
3.3.3 PROPORTIONAL CONTROL 49
3.3.3.1 Initialisation of Parameters: 49
3.3.3.2 Real-Time Feedback Loop: 50
3.3.3.3 Step Response and Ramp Initialisation Method Test: 50
3.3.3.4 Parameter Variation Test: 50
3.3.3.5 Noise Sensitivity Test: 51
3.3.3.6 Controller Gain and Tuning Test: 51
3.3.3.7 Repetition and Averaging Procedure 51
3.3.4 EXPONENTIAL SMOOTHING FILTER 52
3.3.5 COST FUNCTION AND GAIN SELECTION 53
3.3.5.1 Initial Iterative Evaluation 54
3.3.5.2 Proportional Sweep Analysis 54
3.4 ANALYTICAL AIR RECOVERY APPROACH 55

4 RESEARCH RESULTS AND DISCUSSION 56

4.1 MODEL DEVELOPMENT 56
4.2 REGRESSION PERFORMANCE AND MATRIX CONDITIONING 58
4.2.1 REGRESSION ACCURACY AND COEFFICIENT RELIABILITY 58
4.2.2 MATRIX CONDITIONING AND STABILITY 58
4.3 SYSTEM DEVELOPMENT, TESTING AND RESPONSE 60
4.3.1 SYSTEM RESPONSE WITH PROPORTIONAL CONTROL 64
4.3.2 MATRIX CONDITIONING EFFECTS ON CONTROL PERFORMANCE 66
4.3.3 SYSTEM RESPONSE TO ABRUPT CHANGES AND EXTERNAL DISTURBANCES 66
4.3.4 SYSTEM RESPONSE WITH PROPORTIONAL CONTROL AND NOISE 69
4.3.5 SYSTEM RESPONSE TO DIFFERENT SUPERFICIAL GAS VELOCITY SETTINGS 75
4.3.6 SYSTEM RESPONSE WITH EXPONENTIAL SMOOTHING FILTERS 80
4.4 COST FUNCTION ANALYSIS AND CONTROLLER TUNING 87
4.5 DISTURBANCE REJECTION PERFORMANCE OF THE P CONTROLLER 87
4.5.1 ABRUPT SUPERFICIAL GAS VELOCITY REDUCTION SYSTEM RESPONSE 89
4.5.2 ABRUPT FROTH HEIGHT INCREASE SYSTEM RESPONSE 90
4.5.3 GRADUAL FROTH HEIGHT INCREASE SYSTEM RESPONSE 91
4.5.4 GRADUAL FROTH HEIGHT DECREASE SYSTEM RESPONSE 91

5 CONCLUSION **93**

REFERENCES **96**

APPENDIX A **100**

A1. Model Equations 100
A2. Filter configurations 102
A3. Cost Function Analysis Results 105
A4. Analytical Method 106
A4.1 Analytical Approach Methodology 106
A4.2 Modelling Air Recovery using a Quadratic model 107
A5. Analytical Method Results and Discussion 109

LIST OF FIGURES

Figure 1: Mineral flotation under PAR conditions. 2
Figure 2: Attachment of air bubbles to hydrophobic particles. 6
Figure 3: Mineral flotation process diagram. 7
Figure 4: Froth recovery vs particle size for three different collector concentrations from Rahman et al., (2012). 10
Figure 5: Froth recovery vs particle size for three different superficial gas velocities from Rahman et al., (2012). 10
Figure 6: Air recovery vs superficial gas velocity for particle sizes in the feed from Norori-McCormac et al., (2017). 10
Figure 7: Comparison of feed, tailings and dropback copper grades as a function of particle size at different superficial gas velocities from Rahman et al., (2015). 11
Figure 8: Froth recovery vs particle size at different froth depths from Rahman et al., (2015). 12
Figure 9: Air recovery progression across varying superficial gas velocities from Oosthuizen and Craig, (2019). 18
Figure 10: Grade recovery curve illustrating optimising flotation control from Shean and Cilliers (2011). 20
Figure 11: Relationship between air recovery and air rate in a flotation process from (Hadler et al., 2010). 20
Figure 12: Block diagram representation of a PID controller. 26
Figure 13: Block diagram of a Model Predictive Control system. 27
Figure 14: A block diagram of an Extremum Seeking Control system. 28
Figure 15: ESC simulation results for air recovery and aeration rate from (Wepener et al., 2023). 29
Figure 16.: Perturbation Extremum Seeking Control results from Wepener et al., (2023) 30
Figure 17: Top of froth bubble size as a function of superficial gas velocity and froth adapted from Oosthuizen et. al., (2021). 37
Figure 18: Air recovery as a function of superficial gas velocity and froth height adapted from Oosthuizen et. al., (2021). 38
Figure 19: A block diagram of the proposed proportional control system. 46
Figure 20: Air recovery trend over varying froth heights and superficial gas velocities using the developed air recovery model and A, B, C values. 57
Figure 21: Dynamic system response showing the change in bubble size and air recovery over time. 61
Figure 22: Air recovery as a function of increasing superficial gas velocity at varying froth heights. 62
Figure 23: Air recovery as a function of increasing froth height and superficial gas velocity. 63
Figure 24: Air recovery as a function of increasing froth height and constant superficial gas velocity. 63
Figure 25: Air Recovery response with P Control at steady state. 65
Figure 26: Effect of P control on superficial gas velocity at steady state. 65
Figure 27: Air recovery and superficial gas velocity response to a decrease in froth height from 106 to 100 mm at 200 seconds. 67
Figure 28: Air recovery and superficial gas velocity response to an increase in froth height from 106 to 112 mm at 200 seconds. 68
Figure 29: System response when superficial gas velocity is manually decreased at 100 seconds. 69
Figure 30: Air recovery time series at various noise levels and a proportional gain of 0.5. 71
Figure 31: Superficial gas velocity time series at various noise levels and a proportional gain of 0.5. 71
Figure 32: Air recovery response with ramp initialisation. 75
Figure 33: System response of superficial gas velocity with ramp initialisation. 76
Figure 34: Air recovery trends in response to superficial gas velocity step input. 78
Figure 35: Superficial gas velocity response to a repeating step initialisation. 78
Figure 36: Air recovery response to decreasing superficial gas velocity step sequence. 79
Figure 37: Superficial gas velocity response to a decreasing repeating step initialisation. 79
Figure 38: Air Recovery with filter weights of G_1 = 0.6, G_2 = 0.4 82
Figure 39: Air recovery response with K_p = 0.5 for varying filter gain values. 83
Figure 40: Air recovery response with K_p = 0.8 for different filter gain values. 84
Figure 41: Air recovery response for various filter gain values for Kp = 1.5. 85
Figure 42: System response to an abrupt decrease in the superficial gas velocity at t = 200 s. 90
Figure 43: System response to an abrupt increase in froth height from 106 to 112 mm at t = 200 s. 90

Figure 45: System response to a gradual increase in froth height from 106 to 112 mm at t = 200 s. 91
Figure 44: System response to a gradual decrease in froth height at t = 200 s from 106 to 98 mm. 92

Figure A 1: Comparison of the analytical model results for α and Jg at K = 0.5. 110
Figure A 2:Graphical representation of α and Jg over time for K = 0.5. 110
Figure A 3: α calculated using the Taylor Series and the polynomial approximations against Jg at K = 1 111
Figure A 4: Comparison of α and Jg over time, including both the calculated values and their quadratic approximations at K = 1. 111
Figure A 5: Evolution of air recovery and Jg over time at K = 1. 112
Figure A 6: Disturbed Jg, α and quadratic approximations with a σ of 0.1 with K = 0.5. 112
Figure A 7: Disturbed Jg and α over time with σ = 0.1 and K= 1. 113
Figure A 8: Jg and α with a noise level of; σ = 0.2 and K =0.5. 113
Figure A 9: Disturbed α, Jg and disturbed quadratic approximations with σ of 0.2 with K 1. 113
Figure A 10: Disturbed Jg , α and quadratic approximations with a σ of 0.2 and K= 1. 114

LIST OF TABLES

Table 1: Bubble size, froth height and superficial gas velocity values. 39
Table 2: Air recovery, froth height and superficial gas velocity values. 39
Table 3: Sum of square errors for air recovery model validation. 40
Table 4: Sum of squared errors for bubble size model validation 41
Table 5: Regression statistics 57
Table 6: Effect of regularisation on matrix conditioning for different K_p values. 60
Table 7: Table of parameters used in the simulation for α. and Jg at a constant $hf = 106mm$. 70
Table 8: Performance metrics for α across proportional gains and noise levels 72
Table 9: Performance Metrics for Superficial Gas Velocity Across Proportional Gains and Noise Levels 73
Table 10: Table of parameters for simulation when Jg is initialised through a ramp input. 75
Table 11: Simulation parameters when superficial gas velocity is initialised through a repeating step sequence. 77
Table 12: Parameters used for simulation with varying filter gain configurations. 81
Table 13: Performance Metrics for air recovery at K_p = 0.8. 82
Table 14: Performance Metrics for superficial gas velocity at K_p = 0.8. 83
Table 15: Disturbance rejection test results 88

Table A 1: Summary of key equations from literature. 100
Table A 2: Performance metrics for air recovery at K_p = 0.5 103
Table A 3: Performance metrics for superficial gas velocity at K_p = 0.5 103
Table A 4: Performance metrics for air recovery at K_p = 1.5 103
Table A 5: Performance metrics for superficial gas velocity at K_p = 1.5 104
Table A 6: Sweep cost function analysis results. 105
Table A 7: Parameters used in the analytical solution for various K values. 109

NOMENCLATURE

Term	Description	Units
α	Air recovery	-
α_{ref}	Air recovery reference value	-
K_{BSJ_g}	Flotation rate-constant (superficial gas velocity)	-
K_{BS_λ}	Flotation rate-constant - (average froth residence time)	-
$K_{\alpha BF}$	Flotation rate-constant - (Bubble diameter)	-
$K_{\alpha J_g}$	Flotation rate-constant - (Superficial gas velocity)	-
$h_f, (h)$	Froth height	mm
λ_{air}	Froth residence time	hr
D_{BF}	Mean bubble diameter	cm
$D_{BF_{ref}}$	Froth bubble size reference value	cm
J_g	Superficial gas velocity	mm/s
$J_{g_{opt}}$	Optimum superficial gas velocity	mm/s
J_{g0}	Initial superficial gas velocity	mm/s
ξ	Air in overflowing froth	-
v_f	Velocity of overflowing froth	mm/s
w	Weir lip length	cm
Q_g	Volumetric air flow rate	m^3/h

CHAPTER 1

1 Introduction

This chapter introduces the research by providing an overview of the dissertation, focusing on froth flotation as an important mineral processing technique and the significance of Peak Air Recovery (PAR) in enhancing efficiency. It also outlines the research aim and objectives, setting the foundation for the study.

Processing of minerals using froth flotation has been a prominent technique in the mining industry since the early 20th century (Shean and Cilliers, 2011). It is a selective separation process that facilitates the recovery of valuable minerals from gangue minerals based on their distinct surface characteristics. By manipulating several variables such as reagent dosage, pH, and air flow rate, froth flotation allows for economically feasible processing of complex ores, effectively separating a wide range of minerals (Wills and Napier-Munn, 2006).

Over the years, various designs of flotation cells have been developed, including mechanical cells, column cells, and pneumatic cells, each with unique features and applications (Mondal et al., 2021). The efficiency of these cells is influenced by multiple factors, including cell design parameters, operational settings, and the physicochemical properties of the ore and reagents used (Gharai and Venugopal, 2016).

Advancements in flotation technology have focused on improving energy efficiency, enhancing recovery rates, and optimising operational costs. Innovations include the development of new reagents, improved cell designs and the implementation of advanced process control systems.

In recent years, model-based control methods have been introduced to monitor and manage various factors in the flotation process. Extensive research has investigated the influence of different variables on flotation performance. The air flow rate has emerged as a crucial operating parameter affecting froth stability, recovery, and grade (Hadler and Cilliers, 2009).

Further studies revealed that air recovery, the fraction of air entering the flotation cell that overflows as unburst bubbles, significantly impacts froth stability and reaches a peak under specific conditions (Fuerstenau et al., 2007; Hadler et al., 2010; Neethling and Cilliers, 2008).

Operating under peak air recovery (PAR) conditions has been shown to enhance the recovery of valuable minerals, due to reduced water and gangue mineral recovery, as illustrated in Figure 1 (Shean and Cilliers, 2011). Under PAR conditions, froth stability is improved, enhancing the transportation of hydrophobic particles to the concentrate launder while minimising the entrainment of gangue minerals (Barbian et al., 2006). This improvement leads to enhanced selectivity and recovery, making PAR a critical measurement in assessing flotation cell performance. Froth stability depends on multiple factors, including froth height (h_f), superficial gas velocity (J_g), mineralogy, feed grade, particle size distribution, and reagent dosage (Mondal et al., 2021). In the event of froth instability, which is characterised by frequent bubble breakage, hydrophobic particles may sink back into the slurry, lowering recovery rates.

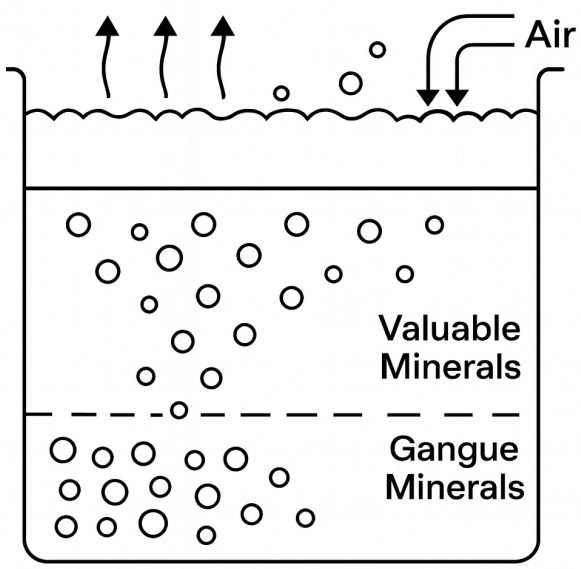


Figure 1: Mineral flotation under PAR conditions.

Traditionally, flotation processes are operated under fixed setpoints for variables such as superficial gas velocity and froth height. However, this approach lacks adaptability to dynamic conditions. Advanced control strategies such as Model Predictive Control (MPC) and Extremum Seeking Control (ESC) have been explored to address this challenge. These

methods, while powerful, often require complex model development or suffer from slow convergence.

In the continuous effort to optimise flotation processes, the role of air recovery is critical as it significantly influences the recovery of valuable minerals. Maintaining operation at the point of Peak Air Recovery has therefore become a key objective in improving flotation control. However, existing approaches, including Extremum Seeking Control, often show slow convergence rates and high computational demands (Ziolkowski et al., 2022, 2021). To address these challenges, the development and evaluation of a proportional control based strategy for maintaining a flotation cell at PAR are discussed in this dissertation. The proposed proportional (P) control strategy, a component of the proportional-integral-derivative (PID) control framework, achieves fast convergence while avoiding the complexity associated with integral and derivative terms.

By critically assessing the benefits and drawbacks of these methods within the scope of optimising air recovery, this study aims to contribute to developing a more efficient and effective flotation control system based on air recovery.

In this context, proportional control offers a simpler and more practical alternative for real time control. This study investigates the use of a proportional controller to manipulate the superficial gas velocity and maintain the flotation system at PAR using regression-based parameter estimation and exponential smoothing filters. The study focuses on evaluating the system's performance under various levels of noise and controller gain as well as the impact of filter gain configurations.

1.1 Problem Statement

Despite the significant advancements in flotation technology and control strategies, challenges remain in effectively controlling flotation cells to operate at PAR. The flotation process involves numerous manipulated variables including superficial gas velocity and froth height, which support the two main control objectives of concentrate grade and recovery.

Existing control strategies may lack the responsiveness and adaptability required to maintain operations at PAR in dynamic process conditions (Hodouin, 2011). Although various

controllers have been developed and implemented, a widely applicable control approach has not been established, as the complex, nonlinear, and disturbance prone nature of flotation systems resists general control design (Shean and Cilliers, 2011; Jovanović and Miljanović, 2015; Quintanilla et al., 2021). Consequently, the primary research question is whether a simple and fast control system can be developed to operate a flotation cell at the point where the air recovery is maximised.

1.2 Research Aim and Objectives

1.2.1 Research Aim

The aim of this research is to design and evaluate a proportional (P) control strategy that effectively maintains flotation cells at PAR by focusing on air recovery optimisation, with the effect on valuable mineral recovery inferred through its correlation with PAR and froth stability.

By critically assessing the benefits and drawbacks of existing control strategies within the scope of optimising air recovery, this research contributes to developing a more efficient and practical flotation control methodology.

1.2.2 Research Objectives

1. Develop a model for air recovery as a function of froth height and superficial gas velocity under both steady-state and dynamic conditions using parameters derived from a literature study.

2. Systematically determine and assess suitable proportional gain (K_p) values to ensure fast convergence, stable operation, and efficient control at PAR.

3. Assess and quantify the performance of the proposed proportional control strategy for PAR, measuring key metrics such as response time and stability to demonstrate its practical applicability in flotation processes.

This work used data digitised from Oosthuizen et al., (2021) to represent a single flotation cell operating under simulated conditions. Dynamic froth and particle interactions beyond this steady state model were not explicitly modelled. This study does not evaluate integral or derivative control terms and is limited to simulation-based testing.

This dissertation is structured to address the challenges of controlling flotation cells at PAR. Chapter 2 reviews existing control strategies, including ESC and PID frameworks. Chapter 3 outlines the method used to develop and validate the air recovery model. Chapter 4 presents results and discusses the P control strategy's performance. While Chapter 5 concludes the research and presents recommendations based on the outcomes.

CHAPTER 2

2 Literature Review

This chapter provides an in-depth review of existing literature on froth flotation processes, emphasising air recovery and its critical role in optimising flotation performance. The review explores the fundamental principles of flotation, the concept of air recovery, the significance of PAR, control strategies in flotation processes, and an analysis of past and current research. This section provides an overview and identifies gaps in the current knowledge that the present study aims to address.

2.1 The Fundamentals of Froth Flotation

Froth flotation is a complex physicochemical process widely used in the mineral processing industry to separate valuable minerals from gangue based on the differences in their surface properties (Barry A. Wills and James A. Finch, 2015). The process involves the addition of reagents to a slurry of finely ground ore, which renders the desired minerals hydrophobic. Air bubbles introduced into the slurry attach to these hydrophobic particles, causing them to rise to the surface and form a froth layer that can be removed as concentrate, as illustrated in Figure 2. The efficiency of this process depends on several operational parameters and the physical and chemical properties of the ore (Kawatra, 2011; Norori-McCormac et al., 2017).

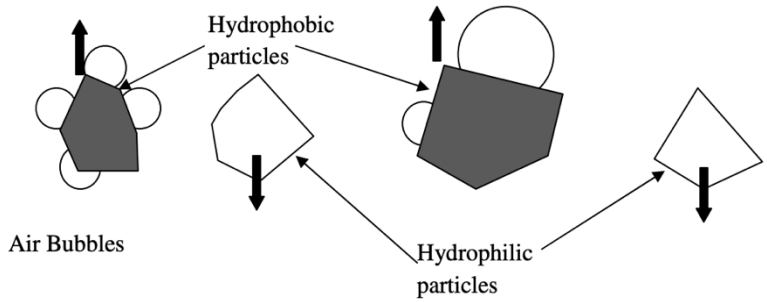


Figure 2: Attachment of air bubbles to hydrophobic particles.

2.1.1 Operational Parameters

Optimising flotation performance requires thoroughly understanding various operational variables, including air flow rate, pH, reagent type and dosage, particle size, froth height and agitation speed. These variables directly influence the efficiency of bubble-particle interactions, froth stability, and overall mineral recovery, especially when aiming to operate at the point of PAR (Hadler and Cilliers, 2009; Shean and Cilliers, 2011). These variables are discussed below, including typical operating conditions for industrial flotation cells to provide context for simulation and control strategies.

Figure 3 illustrates a diagram of a flotation cell, showing its main components. The air supply introduces bubbles into the flotation cell, which interact with the feed slurry. The feed slurry is a mixture of ground ore and reagents while the agitator promotes particle suspension and bubble particle contact. Hydrophobic minerals attach to the bubbles, rising to form a froth layer that overflows into the concentrate launder while the tailings are discarded.

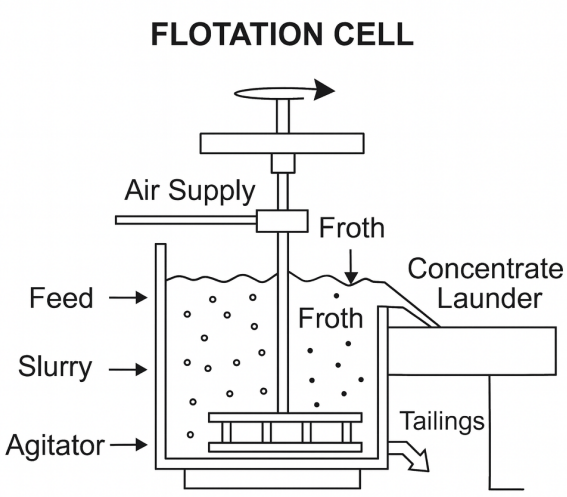


Figure 3: Mineral flotation process diagram.

2.1.1.1 Air Profiles

Air distribution within flotation cells plays a vital role in maximising performance. Various air distribution profiles, such as balanced, increasing, and decreasing, have been analysed to determine their impact on flotation efficiency. Results from the study by Hadler and Cilliers, (2009), suggest that optimising air recovery through tailored air distribution profiles results in higher mineral recovery without compromising the concentrate grade. The article emphasises

the significance of using an air profile that aligns with the PAR rates, thereby optimising the grade and recovery in flotation processes.

Air flow rate, which is often expressed as superficial gas velocity (J_g), is a key parameter affecting froth stability and air recovery. In industrial flotation cells (around 100 - 300 m^3 in size), the superficial gas velocity typically ranges from 1 to 2.5 cm/s, with air flow rates varying from 45 to 150 m^3/h depending on cell size and mineral type (Schwarz and Alexander, 2006; Ostadrahimi et al., 2021).

2.1.1.2 pH Control

pH affects the surface charge of mineral particles, reagent ionisation, and bubble coalescence, thereby influencing mineral hydrophobicity and separation efficiency (Shean and Cilliers, 2011). Proper pH control is essential for achieving selective attachment of valuable minerals to air bubbles while minimising gangue mineral entrainment. The careful adjustment of pH through the use of pH modifiers is critical for optimising the flotation process and obtaining high-grade concentrates of valuable minerals (Barry A. Wills and James A. Finch, 2015; Mondal et al., 2021).

2.1.1.3 Reagent Type and Dosage

The term reagent encompasses collectors, frothers and modifiers where the choice of reagents depends on the mineralogical composition of the ore and the desired selectivity (Mondal et al., 2021; Shean and Cilliers, 2011). Collectors are used to enhance the hydrophobicity of mineral particles, frothers stabilise the froth and control bubble size, and modifiers adjust the pH and depress unwanted minerals. Optimising reagent dosing is important to balance recovery and selectivity. Underdosing may lead to poor attachment of valuable minerals to air bubbles, resulting in lower recovery rates. Meanwhile, overdosing can cause over-stabilisation of the froth, leading to the entrainment of gangue minerals (Barry A. Wills and James A. Finch, 2015). Therefore, carefully optimising reagent types and dosages is necessary to achieve the desired separation efficiency (Oosthuizen et al., 2021).

2.1.1.4 Agitation Speed

Agitation speed, often referred to as impeller speed, is an operational parameter that influences the mixing and dispersion of air bubbles, particles, and flotation reagents within the flotation

cell (Wang et al., 2018). The primary purpose of agitation is to promote the uniform distribution of air bubbles throughout the pulp, increasing the probability of particle-bubble collision (Deglon, 2005).

Another critical aspect is the influence of agitation speed on the hydrodynamics within the flotation cell. Wang et al., (2018) emphasised that flotation equipment must provide suitable hydrodynamics to facilitate interactions between particles and bubbles. The stirring action generated by the impeller not only suspends particles but also creates turbulence, which can enhance the mixing of reagents and improve flotation kinetics. However, the degree of turbulence must be carefully controlled, as excessive turbulence can lead to increased bubble coalescence and reduced froth stability, as discussed by Hadler et al., (2012).

2.1.1.5 Particle Size Distribution

Particle size distribution significantly influences flotation performance by affecting particle-bubble collision probability, attachment efficiency, froth stability, and ultimately air recovery. This relationship governs a range of outcomes from the initial interactions between particles and bubbles to the stabilisation of the froth that enhances effective separation processes.

Fine Particles (<10 µm): Fine particles are shown to achieve high froth recoveries across both low and high superficial gas velocities, highlighting their role in froth flotation processes. Despite their small mass and large surface area, which could potentially reduce collision efficiencies with bubbles, these particles significantly enhance froth recovery. Their ability to remain attached to bubbles, once adhered, reduces their probability of detachment, thereby stabilising the froth. This characteristic not only contributes to reducing bubble coalescence but also increases the rigidity of bubble films. Figure 4 from Rahman et al., (2012) illustrates froth recovery at various particle sizes and at different DDA concentrations, indicating the influence of reagent concentration on fine particle recovery. Additionally, the enhanced recovery of fine particles at varying air rates is demonstrated in the data presented by Rahman et al., (2012) in Figure 5, showing higher air recovery with fine particles under increased air rates.

This finding is supported by Norori-McCormac et al., (2017) who noted how air recovery varies with superficial gas velocities for different particle size distributions, showing that finer particles achieve higher air recovery under increased air rates as illustrated in Figure 6.

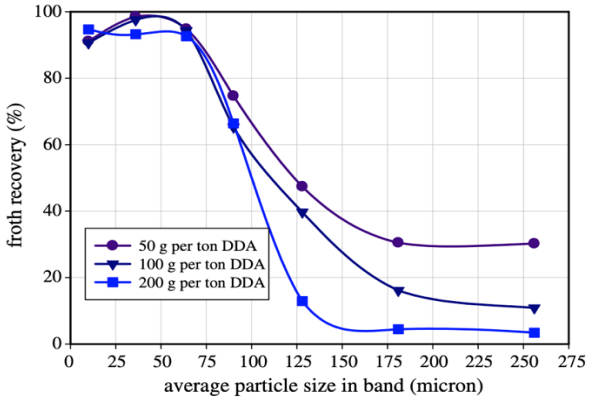


Figure 4: Froth recovery vs particle size for three different collector concentrations from Rahman et al., (2012).

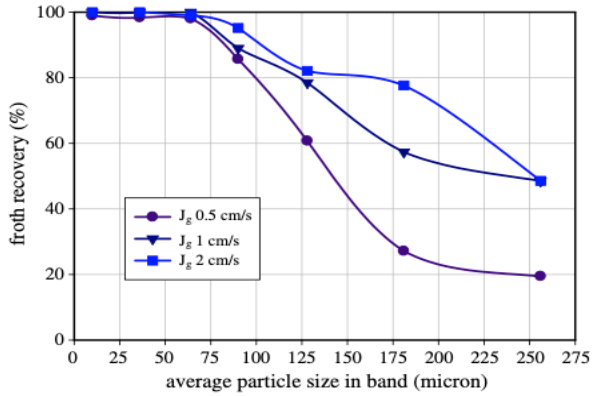


Figure 5: Froth recovery vs particle size for three different superficial gas velocities from Rahman et al., (2012).

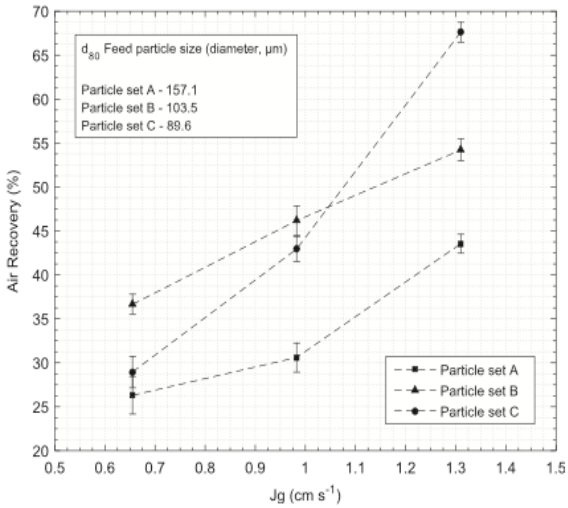


Figure 6: Air recovery vs superficial gas velocity for particle sizes in the feed from Norori-McCormac et al., (2017).

Intermediate Particle Sizes: Particles within the intermediate size range (10 – 150 µm), generally exhibit higher flotation recovery rates due to a favourable balance between collision probability and attachment strength. Rahman et al., (2012) detail that these particles achieve optimal interactions within the flotation cell, which leads to improved recoveries. This size range is key because it often includes particles that are large enough to avoid being carried away by the water in the froth yet small enough to effectively interact with air bubbles.

Additional findings from Rahman et al., (2015) further support the importance of this particle size for achieving peak froth recovery, highlighting their vital role in stabilising froth, which is essential for efficient flotation operation. This understanding is supported by Figures 7 and 8, which presents a comparison of feed, tailings, and dropback copper grades. The graph shows that intermediate particle sizes tend to achieve better grade recovery compared to both finer and coarser sizes, confirming their optimal performance in flotation processes. This optimal range is key not only because of the physical interactions in the flotation cell but also because these sizes are less likely to suffer from mechanical entrainment or to detach under the influence of turbulence, making their recovery more reliable and efficient.

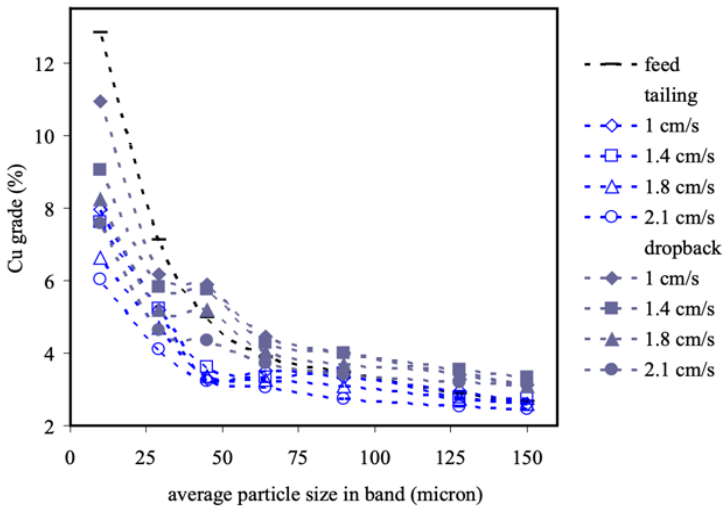


Figure 7: Comparison of feed, tailings and dropback copper grades as a function of particle size at different superficial gas velocities from Rahman et al., (2015).

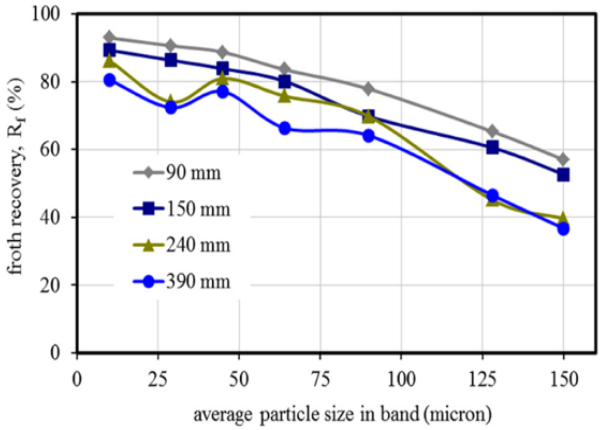


Figure 8: Froth recovery vs particle size at different froth depths from Rahman et al., (2015).

Furthermore, these intermediate-sized particles usually show the highest selectivity in mineral beneficiation, efficiently separating valuable minerals from the gangue. This selectivity is crucial in optimising the economic efficiency of mining operations, reducing the costs associated with further processing and enhancing the overall yield of valuable minerals. Figure 7 shows that maintaining particles within this size range yields a higher concentration of valuable minerals in the concentrate, thus minimising losses to tailings and improving operational efficiency.

Coarse Particles (>150 µm): Coarse particles are more prone to detach from bubbles due to their greater mass and susceptibility to turbulence within the flotation cell. Research by Hadler et al., (2012) indicates that at higher air rates, coarse particles can destabilise the froth by rupturing thin bubble films, thus reducing froth stability particularly when coarse particles are dominant. This behaviour emphasises the importance of optimising air rates to maintain froth stability and maximise recovery.

Managing particle size distribution through controlled grinding, classification, and tailored reagent methods is essential for optimising flotation efficiency. Norori-McCormac et al. (2017) provide important insights into how coarse particles behave under different superficial gas velocities. Their study, demonstrates the performance of coarse particles in flotation processes, showing that these particles respond differently across varying superficial gas velocities, illustrated in Figure 6.

At lower superficial gas velocities, the recovery of coarse particles is typically limited due to insufficient lifting force to raise these heavier particles to the froth layer. As superficial gas velocities increase, there is a noticeable improvement in the recovery of these particles, suggesting that higher air flow rates effectively facilitate the transport of coarse particles to the froth layer. However, the increased turbulence at high superficial gas velocities can cause previously attached particles to detach, leading to increased losses in the tailings. This highlights the balance required in air rate adjustments to optimise the recovery and minimise losses of coarse particles.

Rahman et al., (2012) also emphasise the challenge of stabilising the froth with coarse particles. Their research underlines the need of carefully tailoring flotation conditions to enhance the attachment of these particles to bubbles without causing froth instability.

Additionally, the study by Ziolkowski et al., (2021) shows that by precisely controlling the mill output to target an intermediate particle size range, flotation efficiency can be substantially improved. This is attributed not only to the optimal balance between collision probability and attachment strength, which Rahman et al., (2012) highlight, but also to the enhanced surface properties achieved through specific milling techniques. Such properties are necessary for the effective interaction of particles with reagents and bubbles within the flotation cell, leading to improved froth stability and recovery. Ziolkowski et al., (2021) stress the importance of maintaining the ideal milling conditions to ensure that this intermediate size range can optimise the froth's stability, thus enhancing the overall efficiency of the flotation process. This continuous adjustment and control, facilitated by modern control technologies in grinding operations, are crucial to achieving high recovery rates and consistent flotation performance under varying operational conditions.

2.1.2 Sampling Intervals in Industrial Flotation

To validate models and control strategies, such as those targeting PAR, sampling operations in flotation plants are essential for capturing process dynamics. Sampling intervals depend on the parameter measured and process variability. In copper plants, feed, concentrate, and tailings samples are typically collected over 20 – 60 minutes, across four residence times, to ensure repetitive data (Yianatos et al., 2009). These samples are analysed for particle size, grade, and liberation using sieving and chemical assays.

The effectiveness of flotation control systems is closely tied to the sampling frequency and controller update rate, as these determine how rapidly the control loop can identify deviations and implement corrective action (Shean and Cilliers, 2011). Remes et al., (2007) demonstrated that increasing the assay cycle time led to a decline in control performance, as measured by the integral square error (ISE). Occasional updates can cause delays in the detection of process disturbances, reducing the controller's ability to maintain grade and recovery at optimal set points. Conversely, higher sampling frequencies, done through online elemental analysers such as X-ray fluorescence (XRF), enhance the responsiveness of the control system. However, a trade off exists between sampling frequency and analytical accuracy where shorter measurement times can increase statistical error, while longer times can lead to delayed process feedback (Remes et al., 2007).

Vinnett et al., (2020) performed kinetic tests with collection intervals up to 32 minutes to estimate both the maximum recoveries and the distribution of flotation rate constants. By fitting Gamma distributions to time recovery data, these studies provided insights into particle behaviour across various size fractions. Such analyses inform the selection of sampling intervals for validating process models and control strategies, especially when processes exhibit heterogeneous flotation kinetics.

2.2 Froth Stability

Froth stability is critical in achieving optimal flotation performance, particularly when controlling a flotation cell at its PAR. Once particles are attached to air bubbles, their ability to remain attached until they reach the froth layer at the top of the flotation cell is crucial for effective separation (Barbian et al., 2005). As illustrated in Figure 3, the froth layer must maintain its integrity and stability until it overflows the cell lip. In the event of froth instability, the air bubbles may break, causing hydrophobic particles to sink back into the slurry.

Achieving a stable froth layer is dependent on several factors. Firstly, the mineralogy of the minerals involved significantly influences the attachment of particles to air bubbles. Other factors, such as feed grade and particle size distribution, are vital in determining the flotation performance. The choice and dosage of flotation reagents also impact froth stability and separation efficiency (Mondal et al., 2021; Shean and Cilliers, 2011).

Froth instability, characterised by frequent bubble breakage, causes hydrophobic particles to sink back into the slurry, lowering recovery rates. Studies conducted by Barbian et al., (2003, 2005, 2006), Hadler & Cilliers (2009), and Leiva et al., (2012) investigated and confirmed the relationship between froth stability and air recovery. Their findings demonstrated that concentrate grade and recovery depend on froth stability under various operating parameters. Barbian et al., (2006) further quantified air recovery as a measure of froth stability and overall flotation performance. They found that air recovery peaked at increasing flow rates, which resulted in the highest grade and recovery. However, when air rates exceed this stability threshold, froth destabilisation occurs, causing a reduction in recovery.

One of the most critical factors influencing froth stability is froth height, as it controls residence time, bubble dynamics, and entrainment behaviour in flotation cells. Froth height, defined as the vertical distance between the pulp-froth interface and the overflow lip, directly affects the stability of the froth layer and the efficiency of the separation process. By controlling froth height, operators can manage the time bubbles spend in the froth phase, impacting recovery and selectivity. This makes froth height an important parameter in ensuring consistent flotation performance.

Froth height influences the residence time (λ_{air}) of air bubbles in the froth, a key aspect of maintaining froth stability and ensuring efficient mineral recovery. The equation for λ_{air} is given in Equation 1:

$$\lambda_{air} = \frac{h_f}{J_g} \quad (1)$$

A taller froth height increases the residence time, allowing more time for hydrophobic particles to attach to air bubbles and be transported to the overflow. However, this benefit comes with the risk of excessive bubble coalescence and froth drainage, which can reduce air recovery and compromise the stability of the froth layer (Hadler et al., 2010).

Superficial gas velocity plays a critical role in the interaction between froth height and froth stability. Increasing the superficial gas velocity supplies more air to the system, enhancing bubble-particle attachment and increasing froth height due to the upward force exerted by

bubbles. However, beyond an optimal superficial gas velocity, froth stability diminishes due to the rapid coalescence of bubbles and increased froth drainage (Shean et al., 2017).

Froth height also impacts entrainment, which is the unintended recovery of gangue and water in the concentrate. Taller froths allow more water to drain back into the pulp, reducing entrainment and improving concentrate grade. However, excessively tall froths can delay the upward transport of valuable minerals, particularly in systems with high solid loads, as particles may fail to reach the overflow before bubbles burst (Neethling and Cilliers, 2008). Therefore, optimising froth height is essential for balancing mineral recovery and grade.

The variability of froth height in response to changes in air flow rates, ore feed characteristics, and reagent dosages highlights the importance of including it in flotation control strategies. Sensors capable of real-time froth height measurements, such as laser-based systems, provide critical feedback for control loops designed to stabilise froth height and superficial gas velocity. These systems enable operators to maintain froth stability despite process disturbances, ensuring consistent air recovery and minimising entrainment (Oosthuizen et al., 2021)

Froth height optimisation and effective control strategies ensure the froth layer remains stable, supporting efficient separation and high mineral recovery.

2.3 Air Recovery

The term air recovery was first introduced by Woodburn et al., (1994) and is defined as the fraction of air that is introduced into a pulp that overflows the flotation cell lip as unburst bubbles (Shean et al., 2017; Wei et al., 2014). Air recovery (α) is given by Equation 2.

$$\alpha = \frac{\xi \cdot v_f \cdot h \cdot w}{Q_g} \quad (2)$$

Where:

- α - air recovery,
- ξ - air in overflowing froth,
- v_f - velocity of overflowing froth,

- h - froth height overflowing the weir,
- w - weir lip length,
- Q_g - volumetric air flow rate.

Air recovery serves as an important indicator of froth stability and overall flotation performance. A higher air recovery signifies a more stable froth, which enhances the transportation of hydrophobic particles to the concentrate launder and reduces the likelihood of valuable minerals dropping back into the slurry (Hadler and Cilliers, 2009).

2.3.1 Factors Influencing Air Recovery

Understanding the factors that influence air recovery is important for optimising flotation performance. The primary factors include air flow rate, froth stability, froth height, and operational parameters such as reagent type and dosage.

2.3.1.1 Air Profiles and Air Distribution

Understanding air profiles and air distribution is critical for maximising flotation performance in a bank of flotation cells. An air profile refers to the pattern of air flow rates applied across multiple cells in a flotation bank to achieve a desired metallurgical outcome (Hadler et al., 2010). For instance, an increasing air profile may allocate lower superficial gas velocities to initial cells for higher grade and higher superficial gas velocities to later cells for improved recovery, aligning with PAR conditions. In contrast, air distribution describes the physical dispersion of air bubbles within each cell, ensuring uniform bubble particle interactions to enhance froth stability. Effective air distribution depends on the air flow rate set by the air profile, including cell specific factors like design and agitation speed (Oosthuizen and Craig, 2019). By defining an optimal air profile, operators can control the air distribution across the bank, minimising froth instability and maximising flotation efficiency. Hadler et al., (2010) demonstrated that strategically manipulating air profiles ensures that air distribution supports consistent air recovery, thereby improving overall recovery and grade in the flotation.

2.3.1.2 Froth Height

As previously discussed, froth height is an important parameter that impacts air recovery and flotation performance. An optimal froth height increases the residence time of bubbles in the

froth phase, allowing selective drainage of entrained gangue particles, which enhances concentrate grade while maintaining recovery at the point of peak air recovery (Hadler and Cilliers, 2009). Figure 9 illustrates the air recovery trend at various values of superficial gas velocity where the peak air recovery position is dependent on froth height. These findings align with the work conducted by Hadler and Cilliers, (2009) and Qu et al., (2013), who found that froth height significantly influences air recovery. Increasing froth height up to an optimal value improves air recovery by stabilising the froth and reducing entrainment of unwanted particles. Beyond this optimal height, the benefits diminish, and overall flotation efficiency decreases due to the potential destabilisation of the froth due to bubble coalescence.

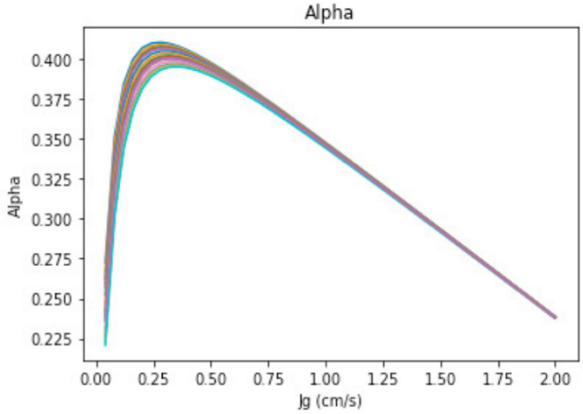


Figure 9: Air recovery progression across varying superficial gas velocities from Oosthuizen and Craig, (2019).

Adding to these observations, Rahman et al., (2015) provide further insights into how froth height affects particle recovery across different sizes, as depicted in Figure 8. This figure shows froth recovery as a function of particle size at different froth height, showing that coarse particles tend to have lower recovery rates, especially at greater froth height. This suggests that as froth height increases, the ability of coarse particles to remain in the froth and contribute to its stability diminishes. The graph indicates that recovery decreases for larger particles as the depth increases, supporting the idea that turbulence and the mechanical limitations of coarse particle flotation become more pronounced at deeper froth levels.

This effect is important as it shows how coarse particles destabilise the froth by detaching from bubbles under turbulent conditions, particularly at higher froth depths where prolonged liquid drainage weakens bubble films, causing bubble coalescence (Hadler et al., 2012). The data in Figure 7 highlights the challenges of managing coarse particles within the flotation cell and

emphasises the importance of optimising operational conditions to minimise these effects, such as adjusting air rates and frother concentrations to better stabilise the froth.

Supporting these findings, Oosthuizen and Craig, (2019) demonstrated through their industrial study that froth height variations are a common and measurable behaviour in flotation systems. Their work introduced the froth stability column, which quantifies froth rise rate and ultimate height, showing that flotation froths typically reach a stable height after growth, influenced by factors such as airflow rate, frother concentration, and froth depth. The study confirmed that higher airflow rates and lower frother concentrations decrease froth stability, while increased froth height can also impact the froth's bursting rate and stability dynamics.

2.4 Peak Air Recovery

PAR is an optimised operational state where air addition is carefully adjusted to enhance flotation efficiency while preserving froth stability. This state is vital as it ensures an optimal air rate that maximises the recovery of valuable minerals without destabilising the froth, which is essential for achieving efficient separation.

Research on PAR has significantly evolved due to the contributions of several key figures in the field. Hadler and Cilliers, (2009) laid the groundwork by linking air flow rates to flotation kinetics and froth stability, highlighting the critical role of air management in optimising flotation cell performance. They demonstrated that controlling air addition could directly affect recovery rates and process efficiency, setting a foundation for future research on optimising these parameters for improved performance and efficiency of flotation cells.

Shean and Cilliers, (2011) contributed to the understanding and modelling of PAR, demonstrating that maintaining an optimal air recovery rate improves the concentrate's grade and recovery. This concept is depicted in Figure 10, which illustrates the trade-off between grade and recovery and highlights the importance of reaching optimum operating conditions for maximum efficiency.

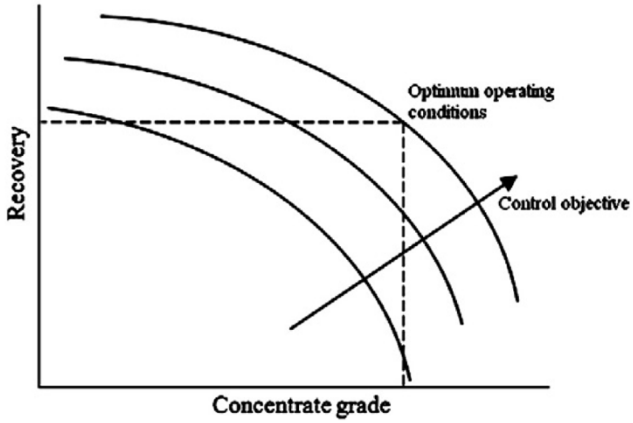


Figure 10: Grade recovery curve illustrating optimising flotation control from Shean and Cilliers (2011).

Hadler et al., (2010) showed that operators can optimise air recovery to achieve PAR by adjusting superficial gas velocity and froth height. As shown in Figure 11, achieving PAR requires tuning superficial gas velocity and froth height to maximise air recovery, ensuring high recovery while maintaining concentrate grade and froth stability. This control strategy requires balancing superficial gas velocity to increase bubble particle transport without destabilising the froth and adjusting froth height to optimise froth residence time for selective gangue drainage (Hadler et al., 2010).

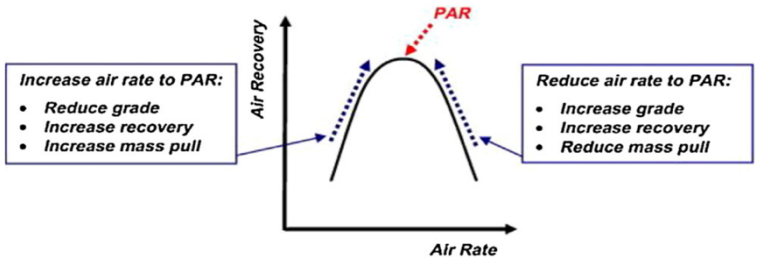


Figure 11: Relationship between air recovery and air rate in a flotation process from (Hadler et al., 2010).

Oosthuizen et al., (2021) developed dynamic models for PAR, integrating real-time data to adjust operational parameters dynamically, ensuring that flotation cells operate under optimal conditions at all times. Building on this foundation, Oosthuizen et al., (2023) incorporated dynamic modelling techniques and machine learning approaches, enabling continuous real-time optimisation of flotation parameters. Their research demonstrated that integrating real-time data with predictive models allows for improved flotation performance, even when faced with varying feed qualities and operational disturbances. In addition, they addressed the

practical challenges of implementing these models in industrial settings, emphasising the need for reliable instrumentation and effective process control strategies.

There are several factors that affect the ability to achieve and maintain PAR. These include:

- Bubble Size and Distribution: Smaller bubbles (0.5 – 2 mm) increase the surface area per unit air volume, enhancing the probability of bubble particle collisions and attachment in the pulp phase, which improves recovery (Schwarz and Alexander, 2006). However, if not properly managed with appropriate frother dosage or air flow rate, smaller bubbles can contribute to unstable froth. Hadler et al., (2010), showed that top-of-froth bubble size is mainly independent of pulp bubble size but plays a vital role in maintaining the froth's structure and air recovery especially at the point of PAR.

- Slurry Properties: Variations in feed properties, such as particle size distribution and mineralogy, can influence the optimal superficial gas velocity for achieving PAR. Smaller particles are more likely to be entrained, while larger particles require higher superficial gas velocity to remain suspended in the froth (Neethling and Cilliers, 2003; Oosthuizen et al., 2021).

- Cell Design and Operation: The design and operation of the flotation cell directly affects air dispersion, bubble formation, and froth stability (Hadler and Cilliers, 2009). Cell design factors such as cell type, geometry, and sparging system, determine hydrodynamic conditions and bubble size distribution. Operational parameters, including impeller speed, superficial gas velocity, and pulp level control, interact with design features to optimise air recovery (Deglon, 2005). The effectiveness of PAR focused control strategies, such as the feedback control of superficial gas velocity and froth height, depends on cell design. Shean and Cilliers, (2011), demonstrated that mechanical cells with adjustable air valves enable dynamic superficial gas velocity tuning, while column cells rely on wash water and pulp level adjustments, which limit superficial gas velocity control flexibility.

2.5 Control Theory in Flotation Processes

Control theory is crucial in optimising flotation processes to enhance performance and efficiency. Flotation control refers to the systematic management of operational parameters to

achieve desired process outcomes, such as maximising mineral recovery and maintaining product quality. Effective control strategies are essential due to the complex and dynamic nature of the flotation processes, which involve multiple interacting variables and disturbances.

Flotation control systems are designed to stabilise and optimise the performance of flotation cells, addressing challenges such as varying feed characteristics and operational constraints. These systems use various sensors and measures to monitor process variables that were discussed earlier, providing data to make real-time adjustments.

2.5.1 Regression Methods in Control Applications

Regression analysis plays a crucial role in control engineering by modelling system behaviour, estimating unknown parameters, and predicting future system states based on observed data (Nakhaei et al., 2012; Liu and MacGregor, 2008). It provides a statistical framework for designing control strategies that adapt to process variations and uncertainties. Unlike purely physics-based models, regression-based control uses historical and real-time data to develop predictive models that improve system performance and stability.

In control applications, regression is often used in:

- Process monitoring and fault detection
- System identification for model-based control
- State estimation and predictive modelling
- Adaptive control techniques

Different regression techniques, including linear regression, multivariate regression, ridge regression, and nonlinear regression, are applied based on system complexity and available data (Cai et al., 2014).

2.5.1.1 Regression Techniques in Control

Regression techniques are widely applied in control systems to model dynamic behaviour, estimate system states, and enhance process stability. These techniques range from basic linear regression to more advanced methods like ridge regression, partial least squares regression, and Moore–Penrose Pseudoinverse applications (Köme and Yazlik, 2024).

Linear and Multivariate Regression

Linear regression models the relationship between an input variable and an output using a linear function:

$$Y = \beta_0 + \beta_1 X + \epsilon \quad (3)$$

where β_0 and β_1 are coefficients estimated from data, and ϵ represents noise.

Multivariate regression extends this to multiple input variables:

$$Y = \beta_0 + \beta_1 X + \beta_2 X_2 + \cdots + \beta_n X_n \quad (4)$$

This method is useful for modelling complex systems where multiple factors influence system outputs (Nakhaei et al., 2012)

In a study by Nakhaei et al., (2012) multivariate nonlinear regression was applied to predict copper (Cu) and molybdenum (Mo) recoveries. The model incorporated input variables such as froth height, gas holdup, and reagent dosage, demonstrating how regression can estimate metallurgical performance based on process data. The study found that Artificial Neural Network (ANN) models outperformed Multivariate Nonlinear Regression (MNR), but regression remained a reliable baseline approach.

Ridge regression

Ridge regression, also known as Tikhonov regularisation, is a regularisation technique used to improve stability in regression models by addressing multicollinearity issues. This method is particularly effective in scenarios where input variables are highly correlated, making standard regression unstable (Hoerl and Kennard, 1970).

Standard regression models can become ill-conditioned when dealing with highly correlated inputs, leading to unstable parameter estimates. Ridge regression addresses this by adding a penalty term $\lambda \, ||\beta||^2$ to the least squares objective:

$$min_\beta \sum (Y_i - X_i\beta)^2 + \lambda \, ||\beta||^2 \quad (5)$$

where λ is a tuning parameter that controls model complexity (Cai et al., 2014).

Ridge regression has been used across various fields, including active flutter suppression for aerospace structures, where it allows real-time adjustments to control inputs, reducing instability in high-speed air flow conditions. As well as in prosthetic limb control where the integration of ridge regression and proportional control enhanced the natural feel and usability of the prosthetic (Igual and Igual, 2024).

Moore-Penrose Inverse

The Moore–Penrose inverse is an important tool in regression-based control applications, particularly when dealing with singular or non-square matrices. This generalised inverse provides solutions to overdetermined and underdetermined linear systems, making it valuable in control system analysis and signal processing (Katsikis et al., 2011).

One major application of the Moore–Penrose inverse is in state estimation and system identification. In practical control systems, full-rank matrices are not always available, and using the Moore–Penrose inverse helps in obtaining least-squares solutions to system equations (Cui et al., 2023).

2.5.2 Control Strategies.

As mentioned, control strategies involve various approaches aimed at optimising flotation performance by adjusting process variables. These strategies include feedback control, feedforward control, Model Predictive Control (MPC), and Extremum Seeking Control (ESC).

- **Feedback Control:** This is the simplest form of control, where the controller responds to deviations in the controlled variable by adjusting the manipulated variable to reduce the error between the measured value and the desired setpoint. For instance, air recovery (α) can serve as the controlled variable, while the superficial gas velocity (J_g) acts as the manipulated variable used to drive the system to the target value (Shean and Cilliers, 2011).

- **Feedforward Control:** This strategy involves adjusting process inputs based on measured disturbances or changes in process conditions such as variations in feed characteristics before these disturbances affect the controlled variable. For example, variations in feed quality could be anticipated through analysis, allowing reagent dosages to be modified in advance to maintain optimal flotation conditions (Hadler et al., 2010)

- **Model Predictive Control:** MPC is a widely used control strategy across various industries, such as chemical processing and automotive systems, using dynamic models to predict future process states and optimise inputs while adhering to constraints over prediction horizons. In flotation, MPC is versatile, adjusting manipulated variables such as superficial gas velocity, froth height, reagent dosage, or pulp level to regulate controlled variables like concentrate grade, mineral recovery, or throughput (Brooks and Munalula, 2017).

- **Extremum Seeking Control**: ESC is an adaptive, model free control strategy applied in industries like energy systems and robotics, where it optimises performance by perturbing process inputs and using feedback to seek the extremum of a performance metric. In flotation, ESC dynamically adjusts variables, such as superficial gas velocity and froth height to find the optimal operating point, maximising or minimising a specific performance criterion (Wepener et al., 2023).

2.5.3 Proportional Integral Derivative Control

PID is a widely adopted feedback control strategy used across industries such as manufacturing and process control. It combines proportional (K_p), integral (K_i), and derivative (K_d) actions to regulate process variables. The proportional term produces an output proportional to the current error, the integral component eliminates steady state error by integrating the error over time, and the derivative component predicts future errors based on the rate of change (Abdul-Adheem and Riyadh Abdul-Adheem, 2017). In flotation, PID controllers are typically applied to maintain controlled variables such as pulp level, air flow rate, or reagent dosage at their desired set points (Shean and Cilliers, 2011; Jiang, 2024). The overall process performance is determined by how optimally those set points are chosen. The basic structure of a PID controller is shown in Figure 12.

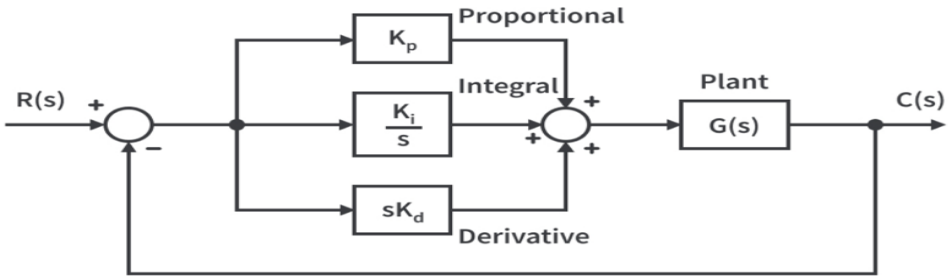


Figure 12: Block diagram representation of a PID controller.

The responsiveness of PID control to changes in system dynamics makes it particularly valuable in flotation processes, where maintaining optimal conditions is critical. Quintanilla et al., (2021) noted that PID effectively stabilises flotation cell variables under minor disturbances. However, traditional PID controllers may struggle in complex systems with non-linear dynamics or changing conditions.

To address these limitations, researchers have explored hybrid PID control systems. Zhang et al., (2017) explored a hybrid control system that combined PID and fuzzy logic to optimise control for flotation column parameters, focusing on optimising aeration rates and pulp levels. This adaptive approach improved recovery rates and concentrated quality by allowing the controller to respond more effectively to varying operational conditions.

A proportional based control system for air recovery optimisation in flotation cells was initially presented in Brooks et al., (2024), a publication that includes results from the early stages of this research. In that study, the controller was derived using the dynamic models developed by Oosthuizen et al., (2021), which describe froth bubble size and air recovery behaviour. The co-authored paper demonstrated that adjusting the superficial gas velocity could drive the system toward operation at the point of Peak Air Recovery.

2.5.4 Model Predictive Control

Another widely used controller is MPC. A basic structure of this controller is shown in Figure 13. MPC is defined as a control strategy that uses an explicit process model to predict the future response of a plant (Bergh and Yianatos, 2011; Quintanilla et al., 2021). Its primary objective

is to optimise future plant performance by determining a sequence of adjustments to the manipulated variables (Jovanović and Miljanović, 2015). As an adaptive controller, it enables the retuning of parameters in response to changing or uncertain process conditions, making it suitable for systems with time-varying dynamics (Jovanović and Miljanović, 2015; Shean and Cilliers, 2011).

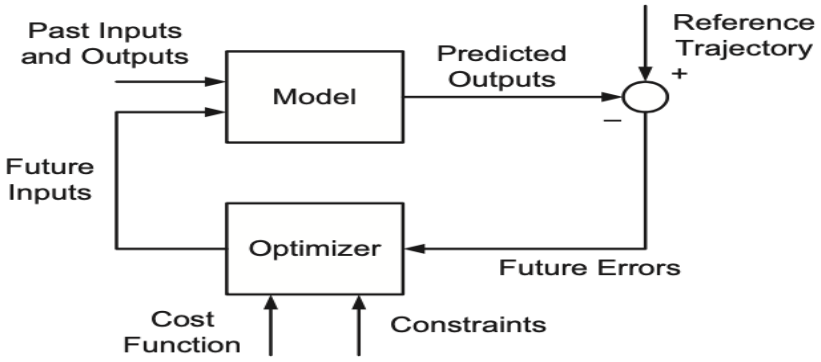


Figure 13: Block diagram of a Model Predictive Control system.

Brooks and Munalula, (2017) implemented a cascaded MPC structure to regulate froth velocity and grade in copper flotation cells. Their approach integrated real-time image analysis for froth velocity measurement and online X-ray fluorescence for grade monitoring, enabling automated adjustments to stabilise flotation conditions. Building on this, Brooks et al., (2019) extended this control strategy to a combined MPC framework that incorporated both milling and flotation operations, improving overall process efficiency by coordinating grinding and flotation control in a unified system.

One of MPC's key strengths is its ability to handle multiple interacting variables and constraints. For example, it can optimise air recovery and froth height, ensuring that the flotation system operates at PAR while maintaining froth stability. Using dynamic models, such as finite impulse response (FIR) or subspace identification techniques, enhances its predictive capabilities, making it well-suited for dynamic and non-linear systems.

However, implementing MPC in flotation presents challenges, including the need for accurate process models, reliable real-time measurements, and significant computational resources.

Despite these challenges, its adaptability and robustness make it an invaluable tool for modern flotation control.

2.5.5 Extremum Seeking Control

ESC is a gradient-based optimisation method that dynamically modifies control inputs based on measured system changes. These modification dynamics are based on the average trend of the gradient, enabling the system to converge toward an optimal operating point without requiring an explicit mathematical model of the process (Ziolkowski et al., 2022, 2021). This adaptive, model-free approach is particularly effective for systems with non-linearities and variability, such as flotation processes. Figure 14 illustrates a block diagram representing a typical ESC control framework showing how the system dynamically adjusts an unknown parameter θ to optimise a process.

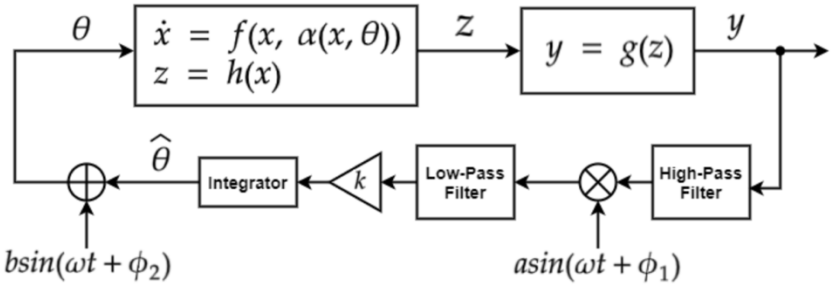


Figure 14: A block diagram of an Extremum Seeking Control system.

In flotation control, ESC has been applied to optimise various parameters influencing the process, such as air flow rate, frother dosage, collector dosage, and pulp level. By continuously seeking the optimal values of these parameters, it assists in maximising the recovery of valuable minerals while minimising the recovery of gangue minerals.

Wepener et al., (2022) demonstrated that ESC can dynamically adjust the superficial gas velocity to maintain flotation systems at PAR despite variations in feed grade or operating conditions. In this study, ESC was implemented by introducing small, continuous perturbations to the aeration rate, allowing the system to explore and gradually converge on the optimal operating point. This method proved to be robust in simulations where it maintained PAR across different conditions, improving mineral recovery and concentrate quality. The study also

emphasised its adaptability, as it responded to external disturbances without operator intervention.

However, despite its potential, significant disadvantages limit the industrial application of ESC in flotation control. One major concern is its slow convergence rate. Because it relies on small perturbations to estimate the gradient direction, it takes time to reach an optimal point, making it ineffective in scenarios requiring quick adaptation to fluctuating conditions. While it is well-suited for steady state optimisation, it struggles in dynamic environments where disturbances occur frequently.

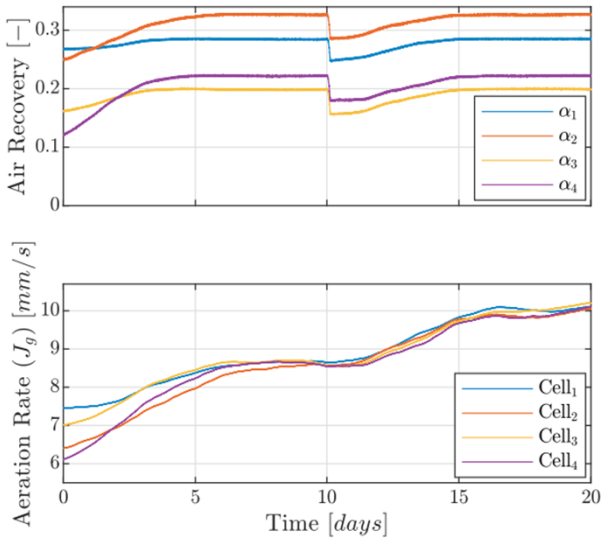


Figure 15: ESC simulation results for air recovery and aeration rate from (Wepener et al., 2023).

This limitation is evident in Figure 15, which presents ESC simulation results for air recovery and superficial gas velocity. The aeration rate shows a gradual increase, requiring approximately 17 days to stabilise close to its optimal value. While this timeframe may seem excessive compared to typical flotation cell residence times, which are often on the order of minutes to hours, the slow convergence highlights a critical challenge of ESC in dynamic environments. The long adjustment period arises from the controller's cautious tuning to avoid destabilising the system. Air recovery exhibits a slow upward trend with fluctuations, including a temporary mid simulation drop, likely caused by disturbances in feed grade or reagent conditions.

Another critical issue is the controller's sensitivity to noise and disturbances, which significantly impacts its performance. The flotation process is naturally noisy due to fluctuations in ore composition, changes in froth stability, and variations in air flow rates. Because it relies on periodic perturbations to explore the input-output relationship, high levels of background noise can mask the true gradient direction, leading to incorrect control adjustments. This problem is further intensified in industrial environments where measurement errors and external process disturbances additionally compromise the reliability of the controller's optimisation process.

Additionally, ESC's reliance on continuous perturbations can destabilise the froth layer, particularly in cases where perturbation amplitudes are not optimally tuned. While ESC's model-free nature allows it to adapt to unknown system dynamics, this comes at the cost of requiring careful tuning of control parameters such as perturbation amplitude and frequency. Incorrect tuning can result in excessive fluctuations in aeration rates, causing the froth structure to collapse or become overly stable, both of which negatively affect mineral recovery.

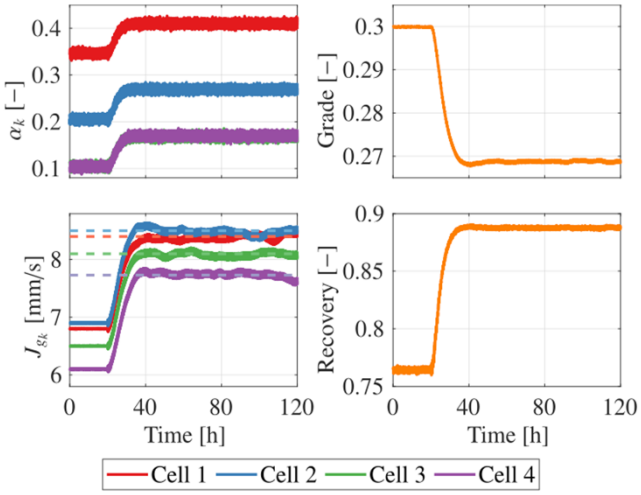


Figure 16.: Perturbation Extremum Seeking Control results from Wepener et al., (2023)

This issue is illustrated in Figure 16, which shows perturbation-based ESC optimisation results show superficial gas velocity. While superficial gas velocity trends toward its optimal value, the adjustment process exhibits non-monotonic oscillations, especially in Cell 3, where the oscillation period varies unpredictably. These oscillations are characteristic to ESC's perturbation mechanism and do not necessarily indicate instability. The grade and recovery

plots further reflect this transient behaviour where grade drops sharply before stabilising at a lower level, while recovery undergoes a temporary adjustment phase before improving.

The challenges associated with ESC are not limited to flotation alone. Similar limitations have been observed in its application to grinding mill optimisation, where it has been used to adjust feed rates and rotational speeds to improve throughput and grind quality. In this context, ESC has been shown to struggle with slow system dynamics, as grinding mills naturally have long residence times and delayed responses to control inputs. As a result, its convergence time is often too slow for practical real-time control, limiting its effectiveness in rapidly changing operational conditions. Additionally, the grinding process involves strong interdependencies between throughput, grind quality, and energy consumption, making it difficult for ESC to simultaneously optimise all variables effectively.

Another disadvantage observed in both flotation and milling applications is ESC's computational demand and implementation complexity. While it eliminates the need for an explicit process model, it requires extensive tuning and parameter selection to function effectively. Finding the optimal perturbation frequency, step size, and adaptation gains is a complex process that varies depending on the system being controlled. Additionally, integrating ESC into existing industrial control frameworks, such as MPC or PID loops, presents additional challenges, as these systems typically require faster response times than ESC can provide.

2.6 Review of Past and Current Research

The use of PAR to assess froth stability, in combination with PID control and MPC, represents a key area of ongoing research in flotation optimisation. Research by Hadler et al., (2012) established that α passes through a peak as aeration increases, indicating that operating at this peak air rate maximises flotation performance. This relationship highlights the importance of maintaining froth stability, which is crucial for achieving optimal recovery and concentrate quality. Oosthuizen et al., (2021) further emphasised that controlling the grade of the flotation product while operating at maximum α is essential for optimising flotation operations.

In their work, Oosthuizen et al., (2021) developed a dynamic flotation model that combines first principles modelling with empirical relationships, incorporating industrial data and froth

imaging to enhance predictive capability. Their model allows key process variables such as air recovery, bubble size, and water recovery to be estimated in real time using only online measurements. This framework supports the application of model-based control strategies and enables the dynamic optimisation of operating conditions, including aeration rate, froth depth, and reagent dosing, to maximise performance under varying plant conditions.

Building on this foundation, Brooks et al., (2024) proposed a simple yet effective control strategy to maintain flotation cells at PAR. Their approach involves estimating the rate of change of air recovery with respect to superficial gas velocity and using this as the input to a proportional controller. The controller adjusts the superficial gas velocity incrementally to drive this gradient to zero, corresponding to the maximum point on the air recovery curve. The estimation of this gradient is achieved through a fitted steady-state model, using parameters derived from measurements of froth depth and air flow. Compared to ESC, their method is significantly faster and requires only a single tuneable parameter. Simulations demonstrated that the controller could effectively steer the system towards PAR even under varying noise levels and changes in froth depth, showing promise for practical implementation in industrial flotation circuits.

However, due to the complex and nonlinear nature of flotation processes, advanced control strategies such as MPC are increasingly being explored. Over the years, Brooks and colleagues have made significant contributions in this field. For instance, Brooks and Munalula, (2017) implemented a cascaded MPC structure to regulate froth velocity and concentrate grade, leveraging real-time image analysis and X-ray fluorescence for predictive control, thereby improving both stability and recovery. Similarly, Quintanilla et al., (2021) developed a dynamic flotation model incorporating froth physics to enable predictive control, demonstrating the potential for improved separation efficiency through advanced modelling techniques. These approaches allow real-time adjustment of operational parameters based on anticipated process behaviour, helping to maintain desired recovery rates and concentrate quality.

More recently, Oosthuizen et al., (2024) advanced the application of MPC by integrating nonlinear control strategies designed to dynamically optimise air recovery. Their method relies on the real-time estimation of process parameters using standard industrial measurements,

allowing the MPC to adaptively follow a grade-recovery curve that maximises efficiency under changing conditions.

The impact of froth stability on flotation performance is further demonstrated in the study by Oosthuizen et al., (2021) which shows that stable froths significantly improve mineral recovery. Continuous monitoring and control of froth stability are therefore critical. The implementation of advanced sensors and online measurement technologies facilitates real-time tracking of froth characteristics, empowering operators to make informed decisions regarding air rates and reagent dosing.

In addition to these developments, optimising flotation banks through recovery profiling has also gained attention. Maldonado et al., (2012) emphasised the value of recovery profiling, suggesting that tailoring the operation of individual cells within a bank can lead to notable improvements in overall recovery and concentrate quality. This aligns with the broader aim of flotation optimisation ensuring that each cell operates under its most favourable conditions.

Air profiles are crucial in maximising flotation performance. The manipulation of air addition across different cells in a bank allows for the optimisation of both the recovery and grade of the product. Strategic control of air profiles can help in balancing the retention time and froth stability across the circuit, enhancing overall flotation efficiency (Oosthuizen and Craig, 2019).

Influential studies by Hadler et al., (2009; 2010; 2012) and Shean et al., (2011; 2017), have further emphasised the importance of achieving optimal air recovery. Norori-McCormac et al., (2017) explored the dynamic relationships between air flow rate, bubble size, and froth stability, providing empirical data that supports the PAR model. Shean et al., (2017) contributed by showcasing practical applications and experimental adjustments in operational flotation cells, highlighting how subtle variations in air rate can significantly affect the overall recovery and quality of the concentrate.

Building on the findings presented in Brooks et al., (2024), this dissertation expands the development of proportional control for air recovery optimisation by broadening the experimental scope and analytical depth. The work extends previous simulations through evaluation of proportional gain settings, noise intensities, and filter configurations. This provided a comprehensive characterisation of controller behaviour. Additional quantitative

performance indicators were introduced to enable comparison of control performance. Furthermore, the effect of different filter configurations on responsiveness and noise suppression was assessed to identify configurations that enhanced stability. Collectively, these extensions address limitations noted in the published study, particularly regarding controller robustness under noisy and transient conditions.

This chapter reviewed the literature on froth flotation, highlighting PAR as a key component for optimising recovery and grade. Existing control strategies, such as ESC exhibit slow convergence while P controllers offer stabilisation within seconds. The review identified a gap for a simple, rapid controller to maintain PAR. Chapter 3 outlines the development of a PAR model and designs a P controller, building on this theoretical foundation.

CHAPTER 3

3 Modelling Method and System Development

This chapter outlines the methodology used to implement and evaluate the proportional control system designed to maintain a flotation cell at the point of PAR. It details the model formulation, simulation setup, and performance assessment procedures used to validate the controller's effectiveness.

Data was extracted from the graphical figures in this publication, and the model was validated and recalibrated through online regression analysis, which was used to iteratively update model parameters based on historical data. The development of the control strategy and simulations were conducted using MATLAB Simulink. All assumptions made during the modelling process are clearly stated, and both the dynamic and steady state properties of the systems are described.

3.1 Modelling Method

The modelling method presented in this section supports the overall aim of the study, which is to develop a control strategy to maintain a flotation cell operation at the point of PAR. The dynamic model developed by Oosthuizen et al., (2021) was used to form the control framework, with data extracted from their published graphs to estimate model parameters through regression analysis. The model was then simplified to a steady state form to enable implementation and evaluation of control strategies for achieving and maintaining PAR.

The dynamic models proposed by Oosthuizen et al., (2021) describe the evolution of froth bubble size (D_B) and air recovery (α) as functions of superficial gas velocity and froth residence time. The governing equations were expressed as:

Froth residence time:

$$\lambda_{\text{air}} = \frac{h_f}{J_g} \qquad (1)$$

Bubble size differential equation

$$\frac{\mathrm{d}}{\mathrm{dt}}D_B = \frac{\left(K_{BS_{Jg}}J_g + K_{BS_\lambda}\lambda_{air} - \left(D_{BF} - D_{BF_{ref}}\right)\right)}{\lambda_{air}} \quad (6)$$

Air recovery differential equation

$$\frac{\mathrm{d}}{\mathrm{dt}}\alpha = \frac{\left(K_{\alpha_{BF}}D_{BF} + K_{\alpha_{Jg}}J_g - \left(\alpha - \alpha_{ref}\right)\right)}{\lambda_{air}} \quad (7)$$

Where:

- $\mathrm{D_B}$: froth bubble size (mm),
- K: fitting parameters,
- J_g: superficial gas velocity (mm/s),
- h_f: froth height (mm),
- $\mathrm{D_{BF_{ref}}}$: froth bubble size reference value (mm),
- λ_{air}: froth residence time (hr),
- α: air recovery,
- α_{ref}: air recovery reference value,

These equations captured the dynamic coupling between bubble formation, gas dispersion, and air recovery behaviour in the froth phase. However, implementing such a detailed dynamic representation in a real time control framework proved to be unnecessary. To enable practical controller design and facilitate faster simulation, the dynamic model was simplified to a steady state form through regression analysis using digitised data from Oosthuizen et al., (2021).

The resulting empirical relationship expressed air recovery as a function of froth height and superficial gas velocity:

Steady state Air Recovery Equation

$$\alpha = \mathrm{A} + \mathrm{B}\frac{h_f}{J_g} + \mathrm{C}J_g \quad (8)$$

Where the coefficients are defined as follows:

- $A = (\alpha_{ref} + K_{\alpha_{BF}} D_{BF_{ref}})$,
- $B = K_{\alpha_{BF}} K_{BS_\lambda} \lambda_{air}$,
- $C = (K_{\alpha_{BF}} K_{BSJ_g} + K_{\alpha J_g})$,

Assumptions:

- The system reaches steady state rapidly, making time derivatives effectively zero.
- The constants $K_{\alpha_{BF}}$, K_{BS_λ}, K_{BSJ_g}, $K_{\alpha J_g}$, are assumed constant over the operating range.

3.1.1 Data Extraction

Data was extracted from a published article by Oosthuizen et al., (2021). This paper did not list the values for the variables of gas velocity, froth height, and the bubble size. Instead, these values were inferred from two graphical representations, Figures 17 and 18, which were used to extract the necessary data.

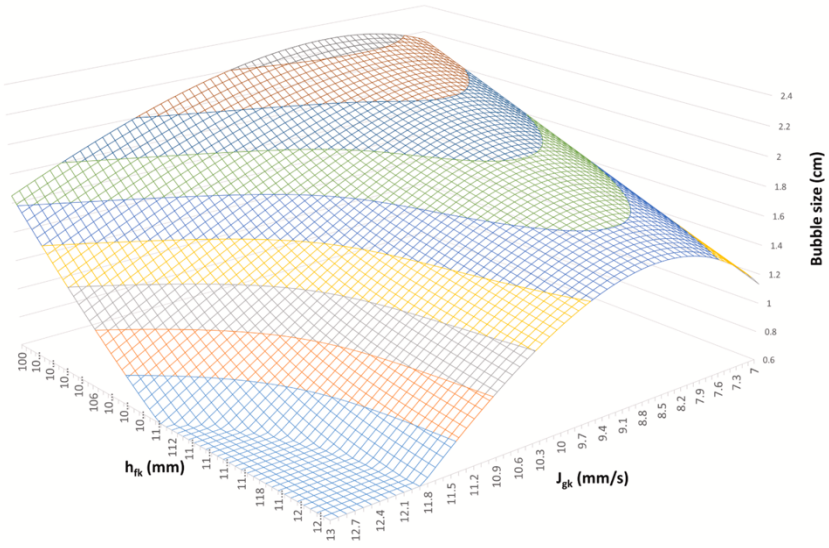


Figure 17: Top of froth bubble size as a function of superficial gas velocity and froth adapted from Oosthuizen et. al., (2021).

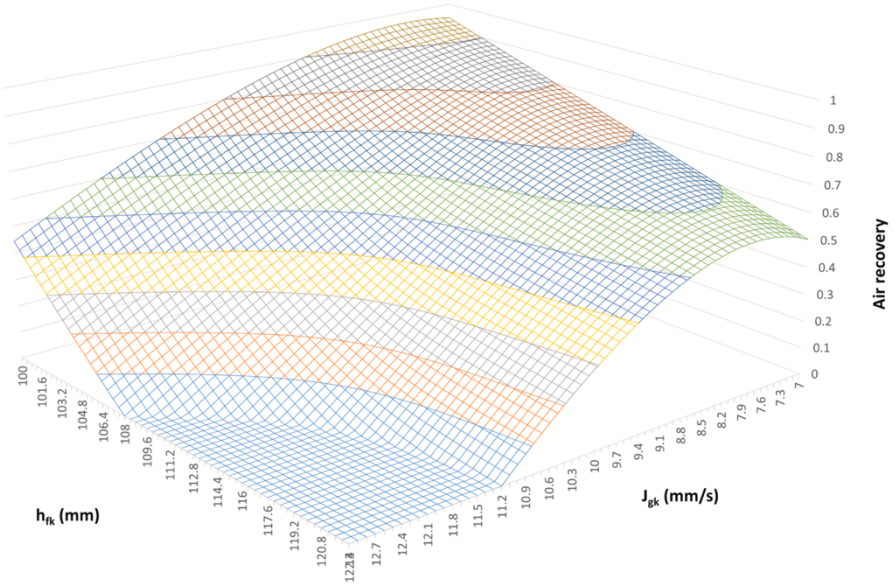


Figure 18: Air recovery as a function of superficial gas velocity and froth height adapted from Oosthuizen et. al., (2021).

The process of extracting numerical values from graphical plots followed a systematic approach. This was accomplished through the following steps:

- Data Point Identification: Data points corresponding to the variables of interest; froth height, superficial gas velocity, bubble size, air recovery, were identified on the graphs. The resolution of the graphs was enhanced using Adobe Illustrator to assist in the reading of data points, with a ruler and pencil to mark and identify the points.

- Recording: The identified values of froth height and superficial gas velocity for both the air recovery and bubble size were entered into Excel spreadsheets. Arbitrary initial values for parameters A, B, C (constrained to be less than 5 and greater than -2) were assigned. A total number of 14 points were extracted for both bubble size and air recovery, as shown in Tables 1 and 2.

Coefficient values assigned:

- A = 5,
- B = 1,
- C = -1,

Table 1: Bubble size, froth height and superficial gas velocity values.

Bubble size (D_B)		
h_f	J_g	D_B
100.60	8.50	2.20
101.80	11.20	1.80
106.00	10.60	1.60
106.40	12.10	1.20
112.00	12.40	0.70
112.00	10.90	1.20
114.40	11.50	0.90
117.40	9.10	1.60
118.00	12.70	0.60
121.00	7.60	1.40
121.60	13.00	0.60
121.60	11.80	0.60
121.60	10.30	1.20
121.60	7.00	1.10

Table 2: Air recovery, froth height and superficial gas velocity values.

Air recovery (α)		
h_f	J_g	α
100.00	13.00	0.45
100.00	7.00	0.80
100.40	7.40	0.95
101.60	11.90	0.65
101.60	10.30	0.45
102.80	8.30	0.83
108.00	12.50	0.09
108.40	10.75	0.33
111.20	9.10	0.56
116.40	11.40	0.01
117.20	10.00	0.32
122.00	8.35	0.51
122.00	11.10	0.04
122.40	7.20	0.54

These variables were used to construct a three-dimensional graph depicting the behaviour of the froth flotation process under varying conditions, capturing the dynamics of the system.

Assumptions in data extraction:

- The scales and proportions in the graphs are accurate.

- Measurement units are consistent across all data points.

3.1.2 Model Validation

Oosthuizen et al., (2021) developed a dynamic model for a flotation cell to predict two performance indicators: the grade and the recovery of the flotation process. The model is based on the differential equations, Equations 7 and 8. The extracted variables were used to calculate the air recovery based on the simplified model expressed by Equation 8. The same procedure was applied to calculate bubble size. The calculated air recovery and bubble size values were compared to their respective values read from the graph, and the errors were computed for all data points. The Sum of Squared Errors (SSE) was then calculated to quantify the overall discrepancy. Excel's Solver tool was used to minimise the SSE by adjusting the values of A, B, and C.

Table 3: Sum of square errors for air recovery model validation.

Air recovery (α)				
h_f	J_g	α (from graph)	α calculated	SSE
100.00	13.00	0.50	0.20	0.30
100.00	7.00	0.80	0.90	-0.10
100.40	7.40	1.00	0.90	0.10
101.60	11.90	0.70	0.30	0.30
101.60	10.30	0.50	0.60	-0.10
102.80	8.30	0.80	0.80	0.00
108.00	12.50	0.10	0.20	-0.10
108.40	10.80	0.30	0.40	-0.10
111.20	9.10	0.60	0.60	0.00
116.40	11.40	0.00	0.20	-0.20
117.20	10.00	0.30	0.40	-0.10
122.00	8.40	0.50	0.40	0.10
122.00	11.10	0.00	0.20	-0.20
122.40	7.20	0.50	0.40	0.10

Table 4: Sum of squared errors for bubble size model validation

Bubble size (D_B)				
h_f	J_g	D_B (from graph)	D_B calculated	SSE
100.60	8.50	2.20	2.20	0.00
101.80	11.20	1.80	1.60	0.20
106.00	10.60	1.60	1.60	0.00
106.40	12.10	1.20	1.10	0.10
112.00	12.40	0.70	0.90	-0.20
112.00	10.90	1.20	1.30	-0.10
114.40	11.50	0.90	1.10	-0.20
117.40	9.10	1.60	1.50	0.10
118.00	12.70	0.60	0.60	0.00
121.00	7.60	1.40	1.30	0.10
121.60	13.00	0.60	0.40	0.20
121.60	11.80	0.60	0.80	-0.20
121.60	10.30	1.20	1.20	0.10
121.60	7.00	1.10	1.20	-0.10

Calculated coefficient values:

A = 4.92,

B = -0.15,

C = -0.28,

3.1.2.1 Dynamic Model

The dynamic model, developed in Simulink, was constructed using a fixed froth height of 106 mm, with superficial gas velocity as the manipulated variable, increased by 0.5 mm/s from an initial value of 5 mm/s in each iteration. The system exhibited fast dynamics, achieving PAR in 20 seconds. Because of the fast system dynamics, all the simulations in this study were conducted under steady state conditions. This approach mirrors standard industrial practice, where steady state conditions are prioritised for process control evaluation to ensure stable, consistent, and efficient long-term process performance by minimising transient effects and enabling reliable optimisation of control parameters.

3.1.2.2 Steady State Model

To model air recovery at steady state, the dynamic equations, Equations 6 and 7, were simplified to a single equation, Equation 8. A detailed justification for this choice is provided

in the results section, where the system's dynamic responses are analysed, and the full derivation of the equation is included in Appendix A.

3.2 Parameter Estimation and Regression Analysis

A multiple linear regression approach was implemented to estimate the coefficients A, B, and C in the air recovery equation to develop a model for air recovery prediction. The regression analysis was performed using the dataset summarised in Tables 1 and 2, where the dependent variable was the measured air recovery, and the independent variables included the froth height to gas velocity ratio $\frac{h_f}{J_g}$ and the superficial gas velocity.

3.2.1 Regression Method

The estimation of these coefficients was achieved through least-squares regression, which aimed to minimise the sum of squared errors between the observed and predicted values of air recovery. The regression process followed the standard matrix formulation:

$$V = (X^T X)^{-1} X^T Y \quad (9)$$

where:

- X represented the matrix of independent variables,
- Y is the vector observed air recovery values,
- V contained the estimated coefficients A, B, and C.

The regression analysis was conducted using Microsoft Excel's Regression Analysis Tool, with R-squared values used to assess how well the model fits the data. The minimum number of data points required for regression analysis depend on the number of independent variables in the model. In a multiple linear regression framework, the general rule is that the number of data points (n) must be greater than the estimated parameters to ensure that the system of equations remains solvable. Specifically, for a regression model with k independent variables, at least k + 1 data points are necessary to avoid an underdetermined system.

The regression equation applied in this study was:

$$\alpha = \mathrm{A} + \mathrm{B}\frac{h_f}{J_g} + \mathrm{C}J_g + \epsilon \quad (10)$$

where A, B, and C represented the estimated parameters. Since the model included two independent variables ($\frac{h_f}{J_g}$ and J_g) alongside an intercept term (A), the theoretical minimum number of data points required for a valid regression was:

$$\mathrm{n_{min}} = \mathrm{k} + 1 = 3 \quad (11)$$

However, using only three data points would not have provided a sufficiently robust estimation, as it would not have allowed for meaningful statistical validation of the regression coefficients. In practice, a larger dataset is beneficial in improving the reliability of coefficient estimates, reducing the risk of overfitting, and improving the model's ability to generalise across different operating conditions. A fixed number of historical data points were incorporated in each regression iteration to enhance numerical stability and reduce the risk of overfitting. A total of 10 data points were used for the regression; however, this number was chosen arbitrarily, without validation, which may have influenced parameter estimation accuracy.

Once the estimated coefficients were obtained, the air recovery equation was used to calculate each data point's predicted air recovery values. This allowed the model to generate expected air recovery outputs based on input variables. In order to assess the accuracy of the model, the residual errors for each data point were calculated using:

$$\epsilon = \alpha_{\mathrm{observed}} - \alpha_{\mathrm{predicted}} \quad (12)$$

where ϵ represents the difference between the observed air recovery values, which were read from the graph, and the predicted air recovery values, with the predicted values calculated using the estimated coefficients and previously specified variables.

After obtaining the estimated constants, the predicted bubble sizes and air recoveries for each data point were calculated using Equations 7 and 8 along with the estimated parameters.

The regression matrix X was constructed as an N x 3 matrix, where N represents the number of data points used in the regression analysis. The first column of X consisted of ones to account for the intercept term A; the second column contained values of $\frac{h_f}{J_g}$, and the third column contained superficial gas velocity values. This formulation allowed the regression to capture both linear and nonlinear dependencies in the system:

$$X = \begin{bmatrix} 1 & \frac{h_{f1}}{J_{g1}} & J_{g1} \\ 1 & \frac{h_{f2}}{J_{g2}} & J_{g2} \\ 1 & \frac{h_{f3}}{J_{g3}} & J_{g3} \\ \vdots & \vdots & \vdots \\ 1 & \frac{h_{fN}}{J_{gN}} & J_{gN} \end{bmatrix}$$

where each row corresponds to an individual observation. The observed air recovery values were stored in the vector:

$$Y = \begin{bmatrix} \alpha_1 \\ \alpha_2 \\ \alpha_3 \\ \vdots \\ \alpha_N \end{bmatrix}$$

Solving for V yielded the coefficient estimates:

$$V = \begin{bmatrix} A \\ B \\ C \end{bmatrix}$$

Additionally, the overall model performance was evaluated using the Root Mean Square Error (RMSE), which quantifies the deviation of predictions from actual values:

$$RMSE = \sqrt{\frac{1}{n}\sum_{n-1}^{n} e_n^2} \qquad (13)$$

The regression model was built on the following assumptions:

- A linear relationship between the independent variables and air recovery,
- The independent variables were not highly collinear,

3.3 System Development

A Proportional (P) control strategy was developed to control the superficial gas velocity to ensure that the flotation cell operated at PAR. The control algorithm utilised real-time data and adapted to changing conditions through online regression analysis. By continuously updating the model parameters and adjusting the superficial gas velocity, the system aimed to operate at the optimal point where air recovery was maximised.

Figure 19 presents the proposed control system designed to optimise flotation cell operations. The first block, representing the plant, defines the flotation cell model using the air recovery equation, Equation 8. It receives inputs of froth height and superficial gas velocity which are used in calculating air recovery, α, using the estimated coefficients. The output of the plant block is the real-time calculated air recovery, which is then used as the input to the regression component. The structure of the control strategy is illustrated in the diagram and outlines the steps required to achieve PAR. Each stage of this process is expanded upon in this section.

The control strategy was initially developed on paper and later implemented in Simulink. The implementation process followed three main stages: parameter identification, control system design, and simulation testing. Each step ensured that the controller was capable of maintaining the flotation system at peak performance with respect to air recovery.

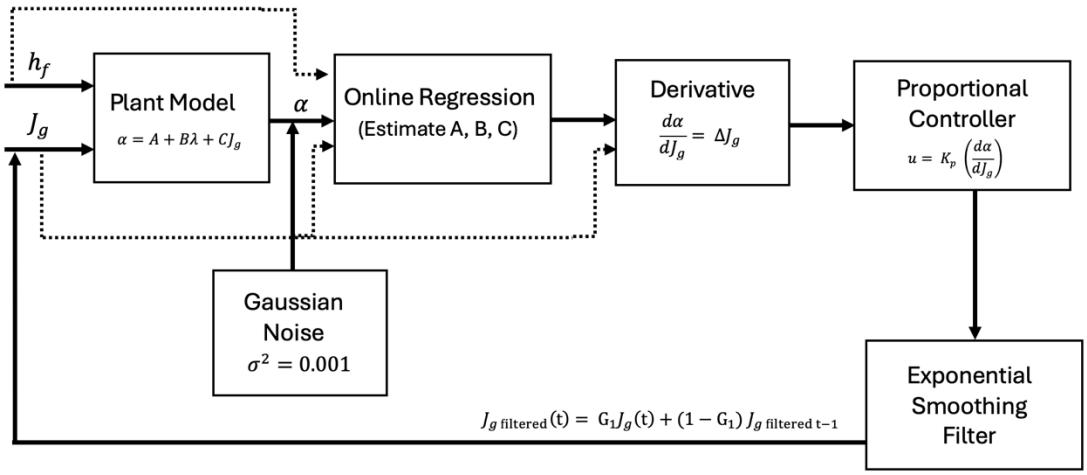


Figure 19: A block diagram of the proposed proportional control system.

In the regression component, the coefficients (A, B, C) were computed using historical data and ongoing performance to ensure the model's precision under steady state operating conditions. The constants are recalculated using the growing number of data points to refine their values, enhancing the model's representation of the process. This recalculation accounts for variations in system behaviour during steady state operation, ensuring reliable parameter estimates. The control system used the derivative of air recovery with respect to superficial gas velocity as the process variable, aiming to reach a setpoint of zero, which corresponds to operating at PAR. To achieve this setpoint, the system incorporates integral action by adding the error manipulated by the gain to the previous output of superficial gas velocity, allowing the proportional controller to effectively minimise steady state error in real time operation using previous data points.

The proportional controller, receiving the derivative output, applies proportional control to adjust the input superficial gas velocity, updating the output incrementally based on Equation 14. This adjustment aims to drive the system towards the desired PAR value, with the incremental update providing an integral like effect to minimise steady state error. An exponential smoothing filter, positioned at the end of the control loop, smooths out the controller outputs to mitigate the impact of abrupt changes or noise in the control signals, ensuring stable control actions are applied to the plant.

$$J_g(t) = J_g(t-1) + K_p e(t) \qquad (14)$$

Initially, tests were conducted without noise to establish a baseline response of the system and assess its natural behaviour under controlled conditions. Once this baseline was established, Gaussian noise was introduced with a mean of 0 and varying standard deviations. The standard deviations were chosen to reflect plausible variations in air recovery, scaled to align with expected steady state fluctuations in similar flotation processes, with various tests conducted at different standard deviations as detailed in the noise sensitivity analysis section (Section 3.3.3.5). This step was crucial for simulating realistic plant conditions where variability and disturbances affect flotation performance. An exponential smoothing filter was applied to mitigate this noise, with its smoothing factor tuned through simulation to balance responsiveness and noise reduction, as evaluated by comparing filtered outputs to the noise free calculated air recovery. While real plant data was unavailable, the filter's effectiveness was validated by analysing its ability to reduce the mean squared error between filtered and noise free air recovery values across different simulation runs at different noise levels, ensuring the control system's stability and performance. By analysing the system's response across different noise intensities, the robustness of the control strategy was assessed, particularly in terms of its ability to achieve and maintain PAR despite external disturbances.

3.3.1 Regression Application in Simulink

The regression method applied to the simulation followed the approach outlined in Section 3.2.1. The algorithm dynamically updated the coefficients by incorporating recent data points in each iteration, allowing the system to continuously improve the model based on the latest operational data and ensuring real-time adaptability. In this simulation, 10 data points were used for each regression iteration. This number was chosen without formal validation, which may have influenced parameter estimation accuracy.

To integrate the regression model within Simulink, the system was designed to perform least-squares regression at each time step, estimating and updating the coefficients of the air recovery equation based on the most recent set of process data. The updated coefficients were then used to calculate predicted air recovery values, which guided adjustments in the control strategy.

A key challenge in implementing regression within Simulink was managing the occurrence of ill-conditioned matrices. To address this, the system calculated the reciprocal condition number (RCOND value) of $X^T X^{-1}$ to measure its stability. If the RCOND value fell below a predefined

threshold (e.g. 1 x 10^{-15}), the matrix was considered ill-conditioned and unsuitable for reliable inversion. In such cases, the system reverted to previously calculated coefficients, preventing instability in parameter estimation. The previous values were then used to reconstruct the matrix, allowing the system to bypass unstable conditions while maintaining operational continuity.

To improve regression stability and mitigate ill-conditioning, the regression model was modified using ridge regression, described in Equation 15, which introduces a regularisation term as a penalty in the least-squares formulation. This regularisation term enhances numerical stability by constraining the magnitude of the fitted coefficients.

$$V = (X^T X + \lambda_{reg} I)^{-1} X^T Y \quad (15)$$

where I is the identity matrix, and λ_{reg} is the regularisation parameter. A value of $\lambda = 0.0001$ was applied to balance model flexibility and stability. This approach prevented overfitting while maintaining numerical stability in the coefficient estimation.

This iterative approach allowed the model to adapt to evolving process conditions while maintaining computational reliability. By continuously updating parameter estimates, the system ensured accurate, real-time predictions of α without being compromised by numerical issues.

3.3.2 Setpoint Calculation Method

The control system's setpoint was defined as the condition where the derivative of air recovery with respect to the superficial gas velocity equals zero. This condition represents the PAR point, the operating point where the flotation system achieves maximum air recovery efficiency.

The derivative of air recovery with respect to superficial gas velocity is given by:

$$\frac{d\alpha}{dJ_g} = -B\left(\frac{h_f}{J_{g^2}}\right) + C \quad (16)$$

3.3.3 Proportional Control

A process model of a flotation cell was developed in MATLAB Simulink to simulate its dynamic behaviour, while a control strategy was implemented to regulate the superficial gas velocity and maintain enhanced air recovery at the peak air recovery point. The simulation environment enabled testing under controlled conditions, such as steady state conditions and varied conditions, such as changes in froth height and abrupt superficial gas velocity adjustments. Gaussian noise was introduced to the calculated air recovery output values to simulate process and measurement disturbances, allowing evaluation of the control strategy's reliability, stability, and adaptability in response to fluctuations in air recovery caused by these disturbances. The goal was to assess the controller's ability to achieve and maintain PAR despite these simulated variations.

The Simulink model was designed to simulate the behaviour of a flotation cell based on a published model, with initial conditions set away from PAR to observe the control strategy's response as it adjusted the superficial gas velocity to reach the optimal operating point for this system. Although the model was not derived from real plant data, it was parameterised using literature based values representative of typical flotation processes, allowing realistic simulation of the system's response to control actions and Gaussian noise induced disturbances. This setup allowed for the observation of the controller's performance under simulated conditions that mimic real world variability.

The following key elements were included in the simulation setup:

3.3.3.1 Initialisation of Parameters:

- Initial condition for superficial gas velocity was set away from its optimal values to test the control system's ability to achieve PAR.

- Parameters A, B, and C were initialised based on regression estimates but were subject to updates every iteration after the first 10 iterations, simulating real-time recalibration and adapting the controller to changing conditions.

3.3.3.2 Real-Time Feedback Loop:

- The control system was implemented as a feedback loop that continuously monitored α and adjusted superficial gas velocity, according to the error signal calculated in the setpoint method.

- Feedback was updated at each time step to simulate a real-time operational environment, allowing the controller to respond dynamically to disturbances and maintain PAR.

A range of tests were conducted in the simulation environment to evaluate different aspects of the control model's performance. Each test was designed to target specific performance criteria, such as setpoint tracking accuracy, response time to reach PAR, stability under Gaussian noise disturbances, and error minimisation in air recovery. These tests provided a comprehensive analysis of the control strategy's reliability and effectiveness in maintaining enhanced flotation performance.

3.3.3.3 Step Response and Ramp Initialisation Method Test:

- A repeating step sequence input was used to observe the system's response when subjected to a sudden change in the superficial gas velocity. The change was introduced from an initial value in both ascending and descending directions. This assessed the controller's ability to stabilise the system around $J_{g_{opt}}$.
- A ramp initialisation method was implemented for superficial gas velocity before the controller was activated. Superficial gas velocity was gradually increased through a ramp input at a constant froth height. Once the ramp phase was completed, the controller took over, adjusting superficial gas velocity to minimise deviations from the setpoint.

3.3.3.4 Parameter Variation Test:

- The variables froth height and superficial gas velocity, were systematically altered to simulate different operational conditions and observe how the controller adapts.

3.3.3.5 Noise Sensitivity Test:

- Gaussian noise was added to α at varying variances (σ^2 = 0.01, 0.001, 0.0001, 0.00001) in separate tests to simulate measurement errors or sensor inaccuracies. Different noise variances were tested to understand their impact on the controller's ability to maintain PAR.

3.3.3.6 Controller Gain and Tuning Test:

- The proportional gain (K_p) was incrementally adjusted to enhance the control system's response time in driving the derivative of air recovery with respect to the superficial gas velocity to its setpoint of zero, corresponding to PAR.

- A cost function, using metrics such as response time to reach PAR, and Gaussian noise disturbance rejection, was used to determine the best proportional gain value.

3.3.3.7 Repetition and Averaging Procedure

To improve statistical reliability of the performance metrics, each disturbance rejection simulation was repeated twenty times under identical input conditions. In each run, Gaussian noise was regenerated using a random seed, ensuring that the disturbance sequence differed across repetitions while maintaining the same noise variance and distribution profile. The random number generator was initialised without a fixed seed, ensuring that each run produced an independent noise sequence.

The Sum of Squared Error (SSE), Standard Deviation (SD), and other relevant control metrics were calculated for each run. Mean values were then computed across all 20 simulations to obtain average performance metrics for each gain and disturbance test. This approach mitigated the effect of outlier runs and provided a more robust assessment of the controller's disturbance rejection capability.

Tests were conducted to evaluate the effect of varying the proportional gain on system performance, including whether high gain values could amplify disturbances or introduce disturbances. This tuning process ensured effective air recovery control while maintaining

system stability and minimising the impact of simulated Gaussian noise disturbances in the flotation process.

3.3.4 Exponential Smoothing Filter

An exponential smoothing filter was applied to the superficial gas velocity in the simulation to reduce the impact of noise and stabilise the system's response. Exponential smoothing is a commonly used technique in signal processing, where recent observations are weighted more heavily than older ones. The weights decay exponentially as data points get older. This approach effectively filters out high-frequency noise while preserving the underlying trend in the data.

The filter operates based on the following formula:

$$s_t = \alpha x_t + (1 - \alpha) s_{t-1}, \quad 0 \leq \alpha \leq 1 \tag{17}$$

where:

- s_t = smoothed value at time t,
- x_t = current data point,
- s_{t-1} = smoothed value from the previous step,
- α = smoothing factor, typically between 0 and 1, determining the weighting between recent and historic data. (Note, in this equation, α represents the smoothing constant and not the air recovery variable used throughout the dissertation).

For this system, the equation was applied to the superficial gas velocity as follows:

$$J_{g\,\text{filtered}}(t) = G_1 J_g(t) + (1 - G_1)\, J_{g\,\text{filtered}\,t-1} \tag{18}$$

where:

- G_1 = filter weight gain for the current J_g value,
- $1 - G_1$ = filter weight applied to the past smoothed J_g value, known as G_2.

The smoothing factor G_1 was adjusted based on the desired outcome between prioritising recent data and historical trends. Configurations where $G_1 > G_2$ prioritised more recent superficial gas velocity values, while $G_1 < G_2$ gave more weight to historical data.

3.3.5 Cost Function and Gain Selection

To evaluate the control system's ability to maintain operation at PAR, a cost function was applied as a comparative performance index. The goal was to identify, assess and rank the performance of proportional gain values under identical conditions, enabling the identification of a suitable operating range.

The cost function is defined as:

$$\frac{d\alpha}{dJ_g} = -B\left(\frac{h_f}{J_{g^2}}\right) + C \qquad (16)$$

$$\frac{d\alpha}{dJ_g} = 0 \qquad (19)$$

The cost function was formulated as:

$$J = \sum_{t=1}^{N} (y_t - r_t)^2 \qquad (20)$$

where:

- J = cumulative squared error,
- y_t = system output α at time t,
- r_t = desired setpoint,
- N = time horizon over which the error is evaluated.

This cost function provided a quantitative measure of control performance, where lower values of J indicated improved alignment between the system's output and the desired setpoint. The squared error metric was chosen because it penalises large deviations more heavily than small

ones, which is relevant because large deviations from PAR can impact valuable mineral recovery and froth stability.

The evaluation process for both the iterative and grid sweep methods used the following test conditions:

- Filter configuration: $G_1 = 0.6$, $G_2 = 0.4$,
- Noise level: $\sigma^2 = 0.001$.

3.3.5.1 Initial Iterative Evaluation

In the preliminary stage, K_p was adjusted iteratively with J, representing the sum of squared errors over the simulation time horizon, calculated after each run to quantify the deviation from the setpoint. If J decreased compared to the previous iteration, the newly selected K_p value was retained, indicating an improvement in system performance. However, if J increased, K_p was revised in the opposite direction. This step was used to establish a range of gains for further evaluation.

3.3.5.2 Proportional Sweep Analysis

In addition to the iterative tuning procedure, a grid sweep was conducted to provide a comprehensive assessment of controller performance across a broader range of proportional gain values. The sweep evaluated K_p values from 0.2 to 1.5 in increments of 0.1, using the cost function defined in Equation 20.

For each gain value, the simulation was executed under identical conditions:

- Disturbance and plant parameters: fixed for all runs.
- Each K_p was simulated 20 times.
- Performance metrics were recorded after each run: SSE, SD, settling time.

The SSE was used as the selection criteria, SD and settling time were used for performance evaluation.

3.4 Analytical Air Recovery Approach

In addition to the simulation-based methodology described in Section 3.3, a simplified analytical model of the system's behaviour near PAR was developed as a supplementary exercise. The objective was to explore the response of the flotation system to proportional type feedback and to provide a theoretical reference for interpreting simulation trends.

The analysis used the air recovery model defined by Equation 8. A proportional constant k was then applied in the governing equation:

$$\frac{\mathrm{d}J_g}{\mathrm{dt}} = \mathrm{K}\frac{\mathrm{d}\alpha}{\mathrm{d}J_g} \tag{21}$$

to examine the effect of feedback strength and sign on convergence to $J_{g_{opt}}$. An empirical and second order quadratic approximation were considered.

Positive K values represented corrective feedback, while negative values were included only to illustrate the instability that results from applying feedback in the wrong direction. This model does not include filtering or process disturbances and was not used to select the final proportional gain. It primarily served as a supplementary validation of the gain and response relationship observed in the study.

Full derivations, equations, and extended methodology for this analytical approach are provided in Appendix A.

CHAPTER 4

4 Research Results and Discussion

This chapter primarily presents the results of air recovery simulations and a single dynamic simulation conducted in MATLAB Simulink, along with a detailed discussion of their implications for optimising flotation performance. In the study by Oosthuizen et al., (2021), graphical representations of air recovery and bubble size models were provided without explicit values for the K constants. By analysing these graphs, as outlined in Section 3.1.1, approximate values were extracted and fitted to the empirical data using the derived air recovery model (Equation 8), with results summarised in Tables 1 – 4.

A simplified analytical model of the system's behaviour near PAR was also developed for theoretical validation and is provided in Appendix A. While informative, this model did not form the basis of gain selection for the final controller configuration.

The discussion in this chapter therefore focuses on the simulation results, which formed the primary basis for gain selection and performance evaluation, while the analytical outcomes serve only as supplementary validation.

4.1 Model Development

To develop the steady state model, the calculated air recovery values were compared against the expected air recovery values, which were extracted from the graphs provided by Oosthuizen et al., (2021) as outlined in section 3.1.1. These expected values served as the basis for determining the empirical parameters A, B, and C through linear regression. This step was crucial for assessing how well the air recovery model replicated the observed air recovery trends across varying superficial gas velocity and froth height values as shown in Figure 20.

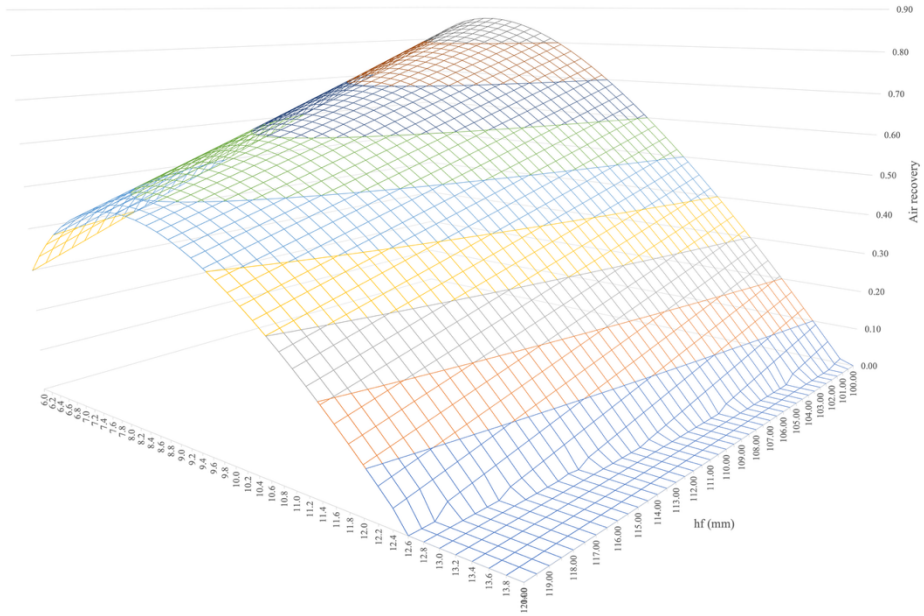


Figure 20: Air recovery trend over varying froth heights and superficial gas velocities using the developed air recovery model and A, B, C values.

The regression analysis was conducted using Excel Solver on the calculated air recovery data values to determine the parameters A, B, C, the regression statistics are presented in Table 5.

Table 5: Regression statistics

Regression statistic	Unit
Multiple R	1
R Square	1
Adjusted R Square	1
Standard Error	1.1763E-16
Observations	6
Intercept	8.84539E-39
X Variable 1	1.18417E-38
X Variable 2	1.65887E-38

The regression analysis results show a multiple R of 1, an R square of 1, and an adjusted R square value of 1, which suggest a perfect correlation and fit between the calculated air recovery values and the extracted data. The standard error of 1.1763E-16 indicates a low residual variability, suggesting that the model effectively captures the data's variability. Additionally, the p-values for the intercept, X variable 1, and X variable 2 were all below p <

0.05, which confirm the model's statistical significance. However, from the data presented in Table 2, the Sum of Square Errors of 0.34 and Root Mean Square Error of 0.16, including a correlation coefficient of 0.69, suggest some variability, which could be due to approximations made when extracting data points from the graph. This indicates that while the model is statistically robust, further refinement of the parameters A, B, and C could improve its accuracy in replicating the data derived from the graph. The points used were selected based on how easily they could be read off the graph. From the regression statistics, it is inferred that the number of data points chosen was satisfactory. This inference aligns with Equation 11 in Section 3.2.1, which defines the theoretical minimum number of data points required for a valid regression as three.

4.2 Regression Performance and Matrix Conditioning

The accuracy and stability of the regression model were key in ensuring reliable parameter estimation for the flotation control system. The regression method followed the approach outlined in Section 3.2.1, where the coefficients A, B, and C were estimated in each iteration using least-squares regression. This process allowed the system to continuously improve its predictions of air recovery based on real-time data.

4.2.1 Regression Accuracy and Coefficient Reliability

The reliability of the regression model was assessed using the coefficient of determination, also known as the R square value (R^2). The regression model produced an R^2 value of 0.74, indicating a reasonable fit between the predicted and observed air recovery values. This suggested that the model captured the general trend of the system.

4.2.2 Matrix Conditioning and Stability

The regression method formed the foundation of the control framework, allowing for the estimation of the coefficients A, B, and C, which were crucial for calculating the air recovery model. This approach enabled real time updates of these parameters, however, several challenges related to numerical stability and reliability were encountered, particularly as the system approached PAR. The effectiveness of the regression method relied on maintaining sufficient variability in the input data, which became more difficult under steady state conditions.

One of the main challenges encountered was the occurrence of ill-conditioning in the regression matrix. As the variability in superficial gas velocity and air recovery decreased near PAR, collinearity developed in the regression inputs, resulting in a nearly singular matrix that produced unreliable coefficient estimates. To quantify matrix stability, the reciprocal condition number (RCOND) of X^TX^{-1} was computed at each iteration. If RCOND dropped below a predefined threshold (i.e. 1×10^{-15}), the matrix was considered ill-conditioned, prompting the system to revert to previously estimated coefficient values. While this prevented sudden, unreliable changes in A, B, and C that could destabilise the controller, it also restricted the system's adaptability by freezing parameter updates, thereby limiting its responsiveness to changing conditions.

To address these issues and reduce the dependence on frozen estimates, ridge regression was implemented by introducing a small regularisation term (λ_{reg}). A λ_{reg} value of 0.0001 was used to prevent overfitting while ensuring that the regression estimates remained stable under varying process conditions. This approach improved the robustness of the parameter estimation process, reducing susceptibility to numerical instability without compromising model accuracy.

The effectiveness of this method is reflected in the results presented in Table 6, which compares the RCOND values with and without regularisation across different K_p values. Without regularisation, the RCOND values were extremely low, demonstrating the poor conditioning of the regression matrix. When regularisation was applied, the RCOND values increased significantly. This confirms that regularisation significantly delayed the onset of ill-conditioning by improving matrix conditioning.

Even though λ_{reg} improved the stability of the regression method, it did not completely eliminate ill-conditioning. The limited variability in superficial gas velocity near PAR continued to affect the system's ability to adapt in steady state conditions. This limitation suggests that additional enhancements would be required to fully resolve these challenges associated with regression in this control strategy. In such cases where ill-conditioning persists, alternative numerical methods such as Singular Value Decomposition (SVD) or the Moore-Penrose pseudoinverse may be employed to further stabilise the regression process and improve the reliability of parameter estimation under poor matrix conditioning.

Table 6: Effect of regularisation on matrix conditioning for different K_p values.

Proportional gain (K_p)	RCOND without regularisation	RCOND with regularisation
0.5	1.94×10^{-16}	6.71×10^{-8}
0.8	2.08×10^{-16}	6.31×10^{-8}
1.5	1.35×10^{-16}	5.20×10^{-8}

To further evaluate the regression method under plant like conditions, noise was introduced into the system. Moderate noise levels proved beneficial, helping maintain sufficient variability in superficial gas velocity and preventing the regression matrix from becoming singular too quickly. However, at higher noise levels, the system struggled to differentiate between actual process dynamics and noise induced fluctuations. This was reflected in the increased variability and unpredictability of the estimated coefficients A, B, and C, which introduced further instability into the control system. These findings highlighted the need to balance noise suppression with the preservation of variability to ensure reliable parameter estimation.

To stabilise the regression process, an exponential smoothing filter was applied to superficial gas velocity, which reduced fluctuations while preserving the variability required for ongoing parameter estimation. The choice of filter configuration had a distinct effect on the system's behaviour. Filters that prioritised recent data improved responsiveness but increased sensitivity to noise, while those that favoured historical data suppressed noise more effectively but reduced the system's adaptability. The optimal filter configuration successfully limited unpredictable coefficient updates without hindering system responsiveness, demonstrating the importance of tuning both the filter and K_p to maintain a stable and adaptive control strategy.

4.3 System Development, Testing and Response

A dynamic simulation of the control system was developed in MATLAB Simulink to evaluate the control system's performance. The simulation was conducted without the addition of noise in the presence of P control. The simulation was run for 200 iterations where 1 iteration = 1 second. Figure 21 shows that air recovery increased rapidly during the first 40 iterations before reaching a consistent level, PAR, while bubble size increased more gradually over time. Both variables approached constant values, with air recovery and bubble size reaching steady state in approximately 80 and 60 iterations respectively.

In this analysis, steady state was defined as the point at which the rate of change of each variable fell below 0.001, corresponding to a variation of less than 2% between iterations. This was adopted with the understanding that a process was considered to have reached steady state once the output variation became negligible relative to the measured range. The fast convergence to steady state supported the decision to focus the remainder of the study on steady state simulations to analyse and enhance the flotation performance. Although the system exhibited dynamic characteristics, this approach ensured computational efficiency by reducing processing demands without sacrificing the accuracy of the process.

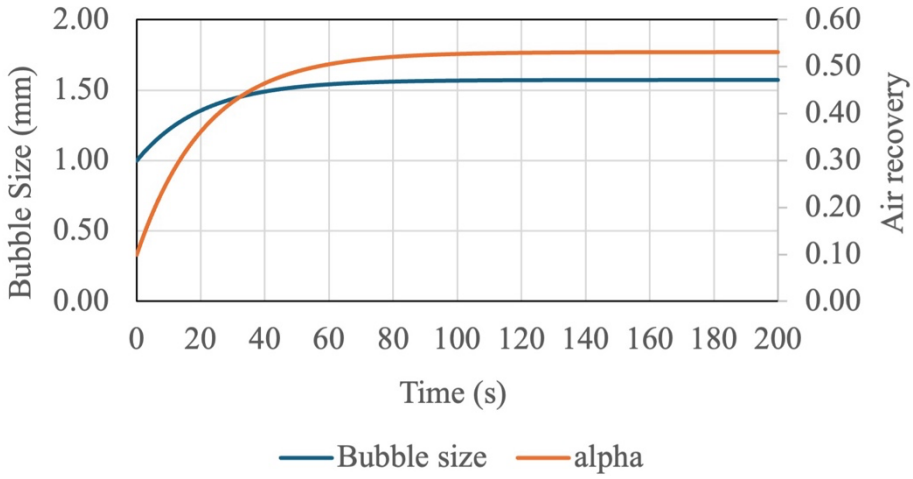


Figure 21: Dynamic system response showing the change in bubble size and air recovery over time.

Following this, a steady state simulation of the control system was developed in MATLAB Simulink. The control system was designed to calculate air recovery as a function of froth height and superficial gas velocity. A P control strategy was implemented to manipulate the superficial gas velocity in order to achieve and maintain operation at PAR. This was important especially when the system was exposed to disturbances such as changes in operating conditions, air flow rate, and feed grade. Based on the model described in Section 3.1.2, the simulation framework incorporated a flotation cell, a control mechanism, and a filtering component, as depicted in Figure 19.

The first steady state simulation was performed without active control or external noise in order to monitor the system's response to changes in superficial gas velocity and froth height under ideal conditions, which provided a baseline for following tests. These simulations also aimed to investigate the effect of froth height on air recovery. Figure 22 illustrates how the PAR line

shifted as the froth height changed. This highlighted how adjustments in froth height influenced system dynamics and the location of PAR. This initial evaluation established the system's fundamental dynamics in the absence of external influences or active control strategies, and the results provided a standard for future controlled tests to be compared against.

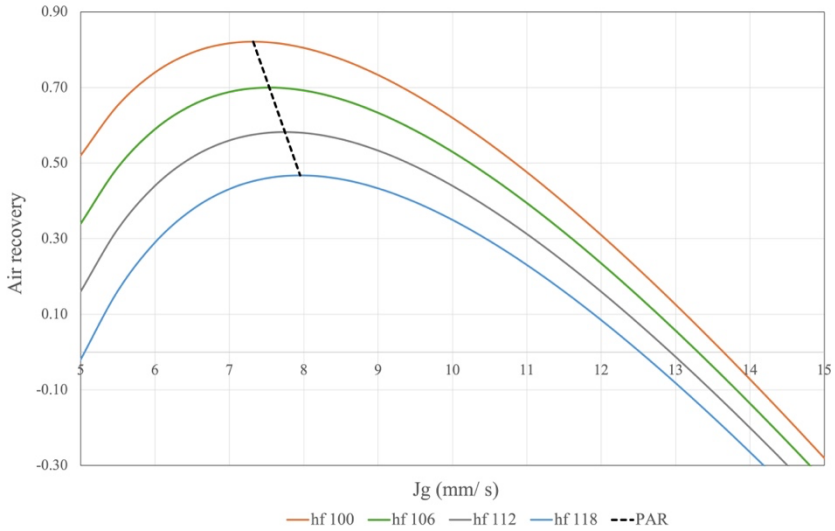


Figure 22: Air recovery as a function of increasing superficial gas velocity at varying froth heights.

Across all froth height values, air recovery increased steadily with superficial gas velocity before reaching a peak and then declining. This trend reflected the expected behaviour of flotation systems, where a peak in air recovery occurred at an intermediate superficial gas velocity. As froth height was varied, the PAR shifted accordingly, demonstrating the system's sensitivity to this parameter. While froth height did not change the overall trend in air recovery, it altered the specific gas velocity at which PAR occurred.

At both low and high values of superficial gas velocity, the model predicted negative values for air recovery, which were not meaningful. As previously discussed, air recovery represents the fraction of air retained in the froth phase and should therefore be positive. These negative values resulted from the quadratic nature of the air recovery equation, which describes a decline in air recovery on either side of the optimal superficial gas velocity. While mathematically consistent with the model's structure, the appearance of negative values highlighted its limitations when applied beyond the expected operating range.

Additional simulations were conducted to test the system's behaviour under varying conditions. One set of simulations increased both froth height and superficial gas velocity, as shown in

Figure 23. Another scenario simulated a constant superficial gas velocity while froth height was increased, as illustrated in Figure 24.

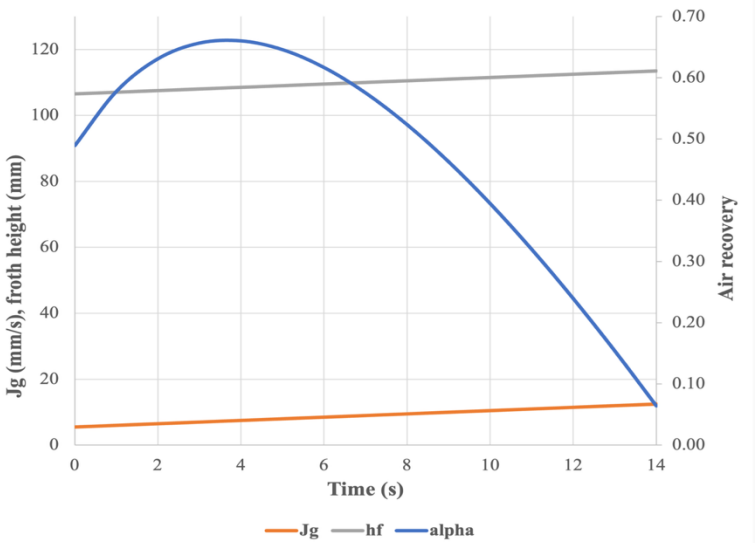


Figure 23: Air recovery as a function of increasing froth height and superficial gas velocity.

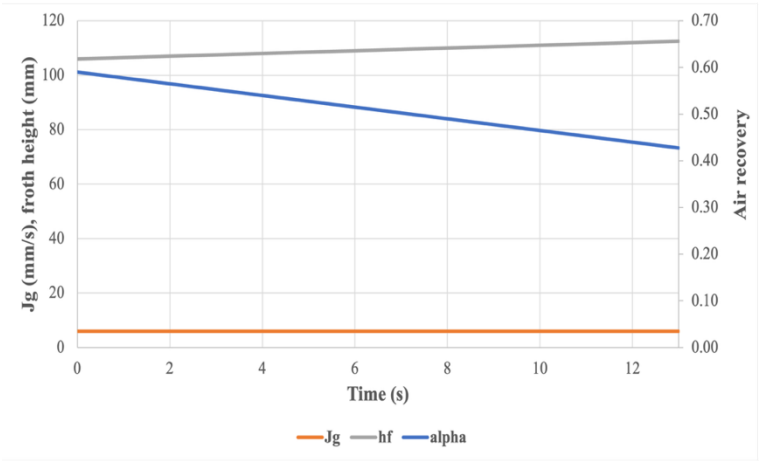


Figure 24: Air recovery as a function of increasing froth height and constant superficial gas velocity.

The results in Figure 23 showed that air recovery varied non-linearly with changes in froth height and superficial gas velocity. As both parameters increased, air recovery initially increased and reached a maximum before gradually declining. This behaviour indicates the presence of an optimal operating range of froth height and superficial gas velocity combinations that promote enhanced air recovery. This behaviour is consistent with previous findings where moderate increases in aeration and froth depth enhance bubble stability and gas dispersion, thereby improving the probability of particle-bubble attachment (Hadler et al., 2010; Shean and Cilliers, 2011). Beyond this optimal point, excessive froth height or superficial gas

velocity may lead to coalescence or bubble collapse, leading to a reduction in air recovery efficiency (Oosthuizen et al., 2021).

In contrast, Figure 24 presented a scenario where superficial gas velocity was held constant and froth height was increased from 106 to 112 mm. This test resulted in a constant downward trend in air recovery without a peak. This steady decrease suggested that the incremental increase in froth height, without an increase in superficial gas velocity, was insufficient to enhance air recovery. The lack of a peak indicated that the fixed gas input could not support the increasing froth height, which potentially led to reduced bubble stability.

The relationship between air recovery and froth is a non-linear relationship, as shown by the varying PAR positions in Figure 22, where different froth heights shifted the point of PAR. This indicated that froth height primarily influenced the location of PAR rather than the overall recovery magnitude. However, the relationship between air recovery and superficial gas velocity was proportional, with superficial gas velocity having played a more dominant role in determining whether a peak was reached. This highlighted superficial gas velocity as the critical control variable that indicated the system's optimal recovery point. To further refine the controller, subsequent simulations were performed with a fixed froth height of 106 mm. This enabled a focused evaluation of how variations in superficial gas velocity affected air recovery.

4.3.1 System Response with Proportional Control

In the following tests, a proportional controller was introduced to manipulate J_g and drive the system to the point of PAR. Following an initial ramp phase where J_g was incrementally increased the controller was initialised. Once the ramp phase was completed, the P controller manipulated J_g in response to deviations from the setpoint, defined in Section 3.1.2 where the derivative of air recovery with respect to superficial gas velocity equals zero.

The impact of different proportional gain values on the system's performance and responsiveness was evaluated to assess convergence behaviour. Higher K_p values, such as 1.5, allowed the system to rapidly respond to deviations, quickly guiding the system towards the targeted air recovery value within fewer iterations. In contrast, lower K_p values, such as 0.5, required more iterations to reach PAR, indicating slower convergence. A moderate K_p of 0.8

demonstrated a balanced response, achieving PAR smoothly without noticeable fluctuations, as shown in Figures 25 and 26. This highlighted the importance of tuning K_p to optimise responsiveness and ensure effective process optimisation. Although the model operates under steady state conditions and the horizontal axis therefore represents iteration steps rather than real time, these results still illustrate the relative speed at which different K_p values drive the system toward the optimal operating point.

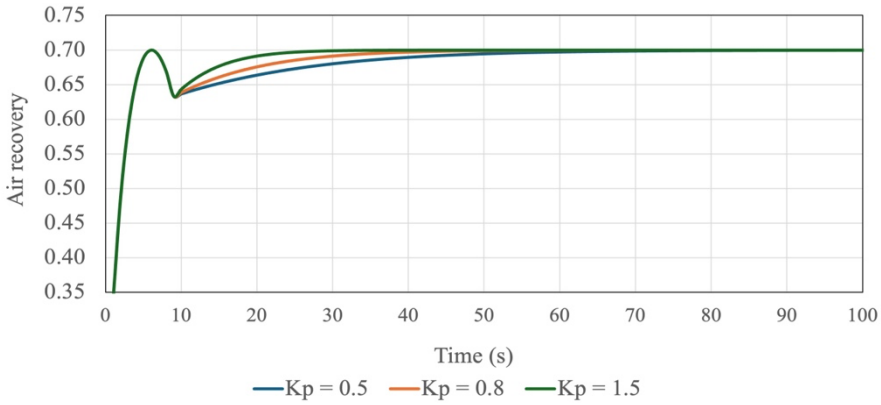


Figure 25: Air Recovery response with P Control at steady state.

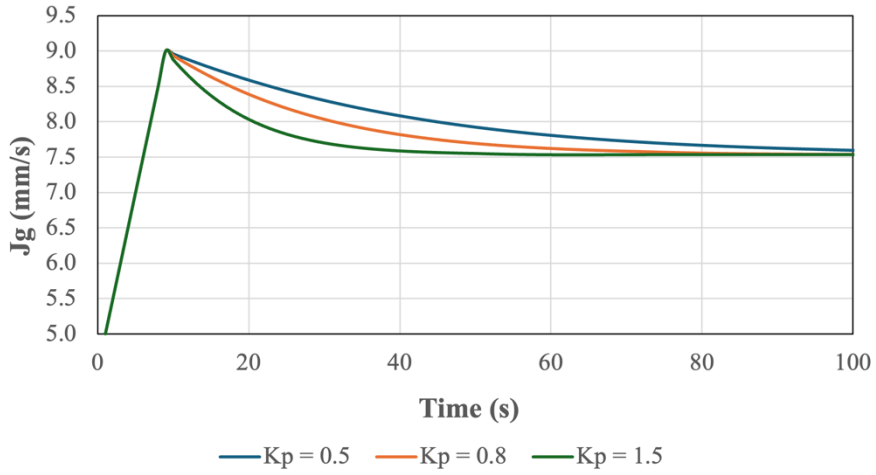


Figure 26: Effect of P control on superficial gas velocity at steady state.

The response trends observed in this section are supported by the analytical results presented in Appendix A, where both implicit and quadratic formulations were derived to describe the relationship between air recovery and superficial gas velocity near the point of PAR. The analytical model confirmed that increasing feedback gain improves the rate of convergence toward PAR. These analytical findings there provide theoretical validation of the simulation results, demonstrating that the proportional control behaviour achieved in MATLAB is consistent with the expected steady state characteristics of the system.

4.3.2 Matrix Conditioning Effects on Control Performance

During the proportional control simulations described in section 4.3.1, inconsistencies were observed in the newly calculated model coefficients A, B, and C. In a number of tests, particularly at higher proportional gains, the estimated constants varied erratically between iterations and often produced unrealistic values. This behaviour indicated a loss of numerical precision within the regression calculations, suggesting that the matrix had become ill-conditioned.

Ill-conditioning occurs when small numerical errors or measurement noise cause large variations in the estimated variables. In this study, the problem arose from near collinearity in the input variables, froth height and superficial gas velocity, which reduced the stability of the regression solution. This was confirmed through RCOND monitoring, which consistently produced very small reciprocal condition numbers ($<10^{-15}$) during unstable simulation runs, confirming that the regression matrix was poorly conditioned. The calculated A, B, and C constants no longer followed a consistent trend, leading to erratic predictions of air recovery.

When ill-conditioning was detected, the system temporarily reverted to using previous stable parameter values until the next iteration produced a valid matrix with acceptable conditioning. This adaptive approach helped maintain simulation stability and prevent propagation of incorrect A, B, and C estimates into the controller. However, while this approach improved short term stability, it was only effective under mild numerical stability and failed to correct more severe ill-conditioning cases. A more advanced corrective method was later introduced in subsequent test cases to improve matrix conditioning and maintain consistent control behaviour. The application and effects of these advanced stabilisation techniques, including exponential smoothing and Tikhonov regularisation are discussed in Section 4.3.6.

Having addressed these numerical instabilities, the next section examines how the proportional control system responded to externally induced changes in froth height and superficial gas velocity, providing insight into the controller's robustness under realistic process variations.

4.3.3 System Response to Abrupt Changes and External Disturbances

Simulations were conducted to evaluate the system's response to abrupt changes and external disturbances, representing real world industrial conditions where fluctuations in operating

parameters occur frequently. The results of these simulations are presented in Figures 27 to 29. These disturbances simulated operational scenarios such as sudden changes in froth height or superficial gas velocity due to upstream fluctuations, equipment variability, or operator adjustments. The system's ability to recover and return to the point of PAR under these disturbances provided insight into its adaptability and control efficiency.

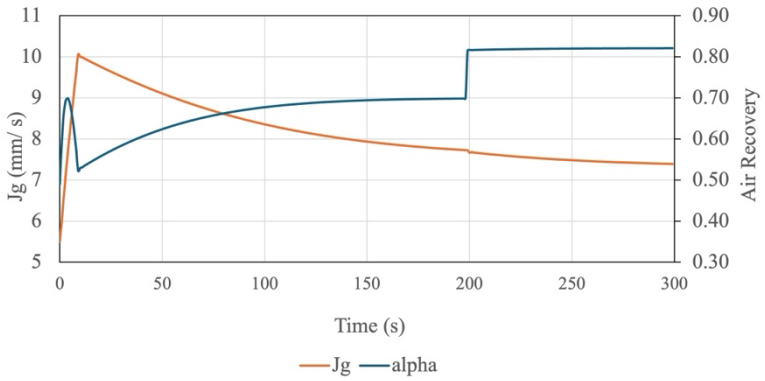


Figure 27: Air recovery and superficial gas velocity response to a decrease in froth height from 106 to 100 mm at 200 seconds.

Figures 27 and 28 illustrate the system's response to froth height variations, a common occurrence in industrial flotation circuits. In Figure 27, froth height was reduced from 106 mm to 100 mm at t = 200 seconds, where 1 second = 1 iteration, representing a step change introduced to assess the controller's ability to re-establish PAR. The decrease in froth height caused an immediate increase in the air recovery value, consistent with the steady state relationship given by Equation 10. This inverse relationship implies that reducing froth height increases air recovery proportionally, producing the step observed in the results.

In contrast, superficial gas velocity showed a small deviation before converging to its expected steady state value of approximately 7.5 mm/s. This demonstrated that the change in froth height did not significantly affect its convergence and that the controller effectively re-established operation at the new PAR point of approximately 0.8, under the modified froth conditions.

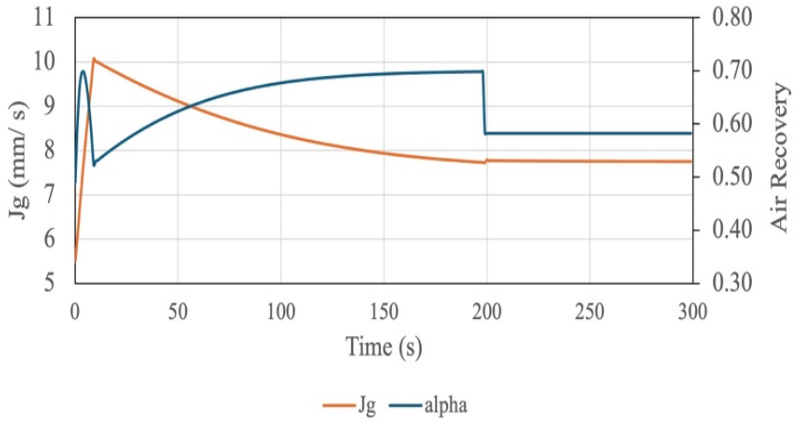


Figure 28: Air recovery and superficial gas velocity response to an increase in froth height from 106 to 112 mm at 200 seconds.

Figure 28 illustrates the system's response to an increase in froth height from 106 mm to 112 mm at t = 200 seconds, where 1 second = 1 iteration. The increase in froth height produced an immediate decrease in the calculated air recovery value. The controller then adjusted the J_g and converged in a few iterations to the new optimum, resulting in a lower PAR point of approximately 0.58. The adjustment in the J_g was minor with a smooth trend, again indicating that the controller re-established operation at the new PAR point.

This shift in the PAR position with varying froth height is consistent with the findings of Hadler and Cilliers, (2009) and Qu et al., (2013), who observed that while optimal froth heights improve air recovery by enhancing froth stability and reducing entrainment, excessive froth heights can lead to froth destabilisation. This occurs due to increased bubble coalescence and reduced residence time of variable particles in the froth, therefore lowering the efficiency of air recovery.

These simulations confirmed that froth height primarily affects the position of PAR, with limited effect on the J_g required to maintain it. The controller consistently re-established operation at the new PAR point following both increases and decreases in froth height, showing smooth convergence and minimal deviation in J_g. These results indicate that the proportional control strategy was robust and capable of handling variations in froth height without compromising flotation performance.

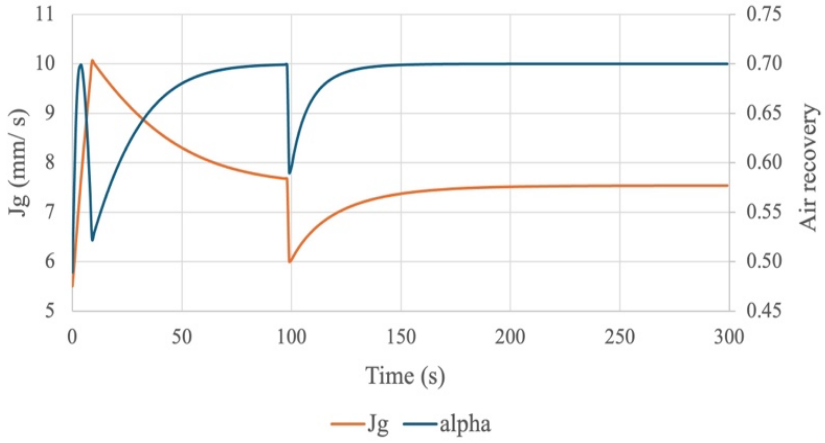


Figure 29: System response when superficial gas velocity is manually decreased at 100 seconds.

Figure 29 presents the system's response to a sudden manual disturbance in J_g introduced at the 100th iteration, simulating unexpected aeration rate fluctuations which can occur during industrial operations. The disturbance resulted in an immediate decrease in α, highlighting its direct dependence on the volume of air entering the froth phase. The controller responded by gradually adjusting J_g through proportional feedback, leading both variables to reconverge to their respective steady state values after a short number of iterations. This behaviour demonstrates the controller's ability to handle abrupt changes in J_g and restore the system to optimal conditions.

Overall, the results suggest that although air recovery is immediately affected by disturbances in superficial gas velocity, the control system effectively compensates for these changes, maintaining stable flotation performance. Additionally, the consistent convergence behaviour observed across different disturbance scenarios indicates a predictable and reliable control response, confirming the robustness of the strategy under the current simulation conditions.

4.3.4 System Response with Proportional Control and Noise

Following the implementation of a P controller to manipulate superficial gas velocity and ensure the system converges at the point of PAR, simulations were conducted with varying levels of Gaussian noise (σ^2 = 0.01, 0.001, 0.0001, and 0.00001). The air recovery and superficial gas velocity responses without noise were used as a benchmark for ideal performance. The controller was initiated after an initial ramp up phase, with responses observed over a 100 iteration period, where 1 iteration = 1 second, with the results illustrated

in Figures 30 and 31. Table 7 contains a list of the parameters used in these simulations and Tables 8 and 9 contain a summary of the results.

The introduction of noise progressively challenged the system's ability to maintain setpoints. At the lowest noise level (σ^2 = 0.00001), the system closely tracked the desired target values (α = 0.7, J_g = 7.5 mm/s), with minimal fluctuations. However, as the noise intensity increased, fluctuations became more pronounced, especially in superficial gas velocity due to its dependency on the air recovery error propagated through the P controller.

To assess the system's performance, five statistical metrics were used:

- MAE (Mean Absolute Error): measures the average deviation from the setpoint.
- RMSE (Root Mean Square Error): highlights larger deviations by penalising higher errors.
- SD (Standard Deviation): quantifies variability around the setpoint.
- IAE (Integral of Absolute Error): assesses the cumulative magnitude of deviations over time.
- CPI (Control Performance Index): provides an aggregate indicator of control quality, where higher values represent better performance.

Table 7: Table of parameters used in the simulation for α. and J_g at a constant h_f = 106mm.

Parameter	Value	Unit
Froth height (h_f)	106	mm
Noise Variance (σ^2)	0.01; 0.001; 0.00001	
Proportional Controller Gain (K_p)	0.5, 0.8, 1.5	

Figures 30 and 31 illustrate the time series behaviour of α and J_g showing the impact of increasing noise levels on system stability and control performance. As the noise level increases from σ^2 = 0.00001 to 0.01, both α and J_g exhibited growing fluctuations around their respective setpoints. The performance of the controller was evaluated across three proportional gain settings, K_p = 0.5, 0.8, and 1.5. The results demonstrated that the proportional gain significantly influenced both the system's responsiveness and its sensitivity to noise.

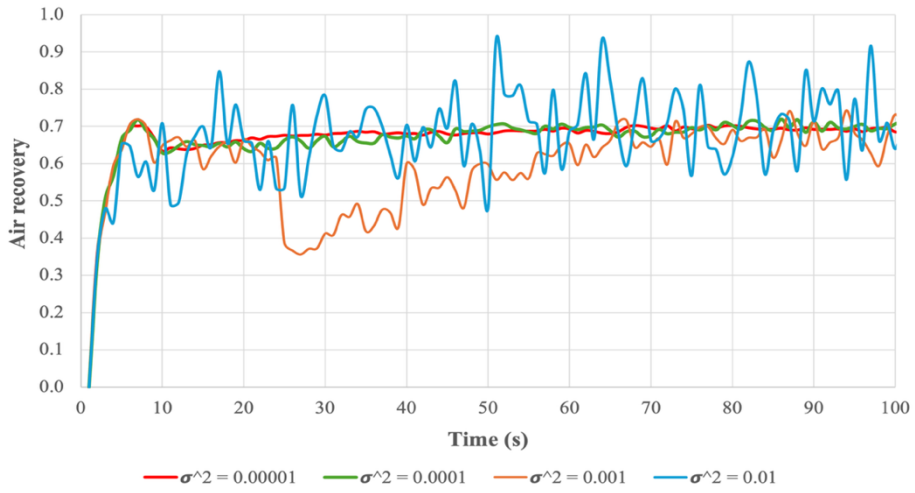


Figure 30: Air recovery time series at various noise levels and a proportional gain of

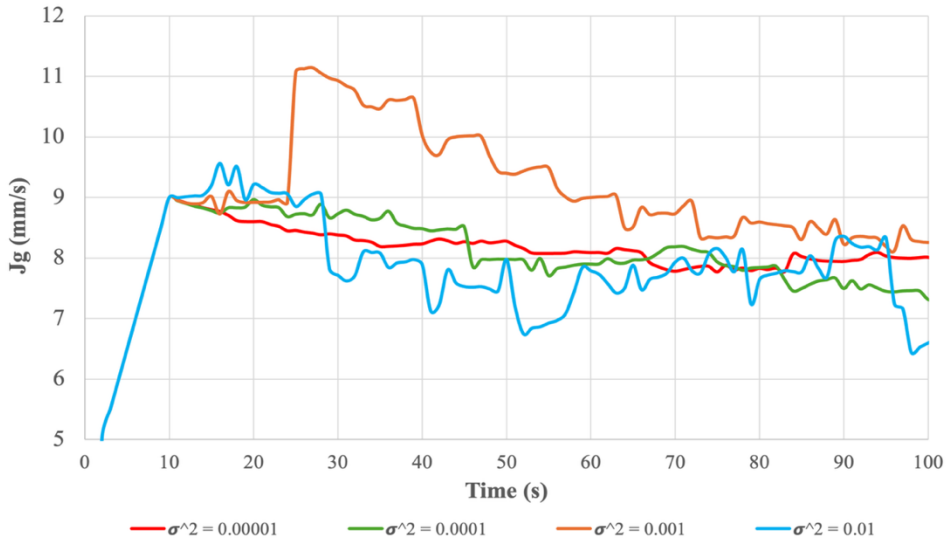


Figure 31: Superficial gas velocity time series at various noise levels and a proportional gain of 0.5.

At K_p = 0.5, the controller demonstrated stable but moderate performance under noise only conditions. When measurement noise was minimal (σ^2 = 0.00001), error metrics such as MAE, RMSE, and SD for air recovery were close to zero, with the CPI reached its highest value, indicating highly stable control. As the noise level increased to a moderate level (σ^2 = 0.0001) the controller maintained low error values and a good CPI, as shown in Table 8. The low IAE values further confirmed that cumulative error remained low across all noise levels, demonstrating effective noise tolerance for air recovery.

The corresponding performance metrics for superficial gas velocity in Table 9 showed a similar trend. MAE and SD values remained low across all noise levels, and CPI increased substantially as noise was reduced, confirming that the controller maintained stable regulation of the superficial gas velocity even under higher measurement noise. Together, these results indicate that the control system provided accurate and reliable steady state regulation of both air recovery and superficial gas velocity, reflecting a strong robustness to measurement noise.

Table 8: Performance metrics for α across proportional gains and noise levels

Proportional gain (Kp)	Noise level (σ^2)	MAE (α)	SD (α)	CPI (α)	RMSE (α)	IAE (α)
0.5	0.01	0.07	0.10	9.41	0.09	7.51
	0.001	0.08	0.09	8.27	0.12	8.55
	0.0001	0.02	0.04	43.26	0.04	1.47
	0.00001	0.01	0.01	121.58	0.01	0.58
0.8	0.01	0.23	0.15	3.09	0.26	22.85
	0.001	0.03	0.03	25.61	0.04	2.76
	0.0001	0.01	0.01	63.12	0.01	1.12
	0.00001	0.01	0.01	80.33	0.01	0.88
1.5	0.01	0.09	0.12	79.02	0.12	9.59
	0.001	0.03	0.05	23.53	0.05	3.00
	0.0001	0.01	0.01	90.14	0.01	0.78
	0.00001	0.00	0.00	261.68	0.00	0.27

Table 9: Performance Metrics for Superficial Gas Velocity Across Proportional Gains and Noise Levels

Proportional gain (J_g mm/s)	Noise level (σ^2)	MAE (J_g mm/s)	SD (J_g mm/s)	CPI (J_g mm/s)	RMSE (J_g mm/s)	IAE (J_g mm/s)
0.5	0.01	0.39	0.50	19.03	0.51	39.81
	0.001	1.08	0.80	6.93	1.34	109.35
	0.0001	0.33	0.70	23.04	0.73	29.62
	0.00001	0.20	0.20	44.16	0.22	17.20
0.8	0.01	2.30	1.25	3.26	2.56	232.20
	0.001	0.33	0.38	22.51	0.45	33.65
	0.0001	0.30	0.22	24.81	0.37	30.53
	0.00001	0.30	0.22	25.29	0.37	29.96
1.5	0.01	0.84	1.03	8.98	1.11	84.37
	0.001	0.28	0.36	27.11	0.37	27.95
	0.0001	0.63	0.48	12.00	0.78	63.13
	0.00001	0.07	0.10	104.37	0.11	7.26

At the highest noise level ($\sigma^2 = 0.01$), the controller's performance at $K_p = 0.5$ declined, as expected. Air recovery showed a slight increase in error, with a slight rise in MAE and RMSE, while the standard deviation remained relatively constant. The IAE increased, and CPI dropped drastically compared to the lowest noise level, indicating reduced control efficiency. Similarly superficial gas velocity exhibited higher error and variability, though the system continued to avoid extreme fluctuations. These results confirm that while K_p performs well under low noise conditions, it is less effective in attenuating higher levels of noise, leading to increased error accumulation and reduced performance.

As shown in Tables 8 and 9, increasing the gain to $K_p = 0.8$ enhanced system responsiveness but also made it more sensitive to noise. At moderate noise ($\sigma^2 = 0.001$), the controller maintained strong performance for both air recovery and superficial gas velocity, although slightly cumulative errors were observed compared to $K_p = 0.5$. This suggests a trade-off between faster convergence and error accumulation.

However, at the highest noise level, performance deteriorated significantly. Air recovery exhibited increased fluctuations and reduced control, while superficial gas velocity showed

increased variability. The controller struggled to reject the high frequency noise, leading to amplified disturbances and reduced control quality.

At K_p = 1.5, the system remained stable under moderate noise, achieving low error and high consistency across all metrics. When exposed to the highest noise level, some degradation in performance was observed, but it remained less severe than at K_p = 0.8. The controller maintained relatively better tracking and convergence, particularly for superficial gas velocity.

At a moderate noise level of $\sigma^2 = 0.001$, all gain settings demonstrated acceptable performance. However, K_p = 0.5 consistently resulted in the lowest error accumulation and the most stable CPI, especially for J_g, where higher gains resulted in greater variability. This stability is important for maintaining reliable and controllable system behaviour under realistic industrial conditions.

Across all gains, RMSE and IAE trends showed that both the magnitude and accumulation of error increased with an increase in noise level, especially for superficial gas velocity. As shown in Figures 30 and 31, the system consistently converged to a steady state operating point within a few iterations under low to moderate noise. At higher noise levels (σ^2 = 0.001 and 0.01), however, fluctuations increased, and the system often failed to reach a clearly defined steady state. Notably, the superficial gas velocity tended to settle at around 8 mm/s, which is higher than the target of 7.5 mm/s, reflecting the controller's response to variations in air recovery introduced by noise.

The metrics reinforced the visual trends, confirming that the performance of both air recovery and superficial gas velocity declined as noise increased, particularly under higher proportional gains.

Based on these observations, a moderate noise level of $\sigma^2 = 0.001$ was selected for subsequent simulations. This noise level introduces a realistic degree of measurement variability commonly encountered in industrial flotation systems while maintaining manageable error levels and stable convergence. It provides a balance between industrial process variability and controller responsiveness, ensuring meaningful evaluation of the control strategy in future tests.

4.3.5 System Response to Different Superficial Gas Velocity Settings

The system's response to the ramp initialisation method and addition of noise was analysed to determine its ability to reach and maintain PAR. In previous tests, J_g was ramped in 0.5 mm/s increments to assess initial system behaviour. The following tests were conducted to investigate the impact of how J_g is initialised on the system and to determine if the position of J_g relative to its optimal value affects system performance. The parameters used are listed in Tables 10 and 11, with results illustrated in Figures 32 to 37.

Table 10: Table of parameters for simulation when J_g is initialised through a ramp input.

Parameter	Value	Unit
Froth height (h_f)	106	mm
Proportional Controller Gain (K_p)	0.5; 0.8; 1.5	
Superficial gas velocity (J_g) ramp	1	mm/s

The results shown in Figure 32 indicate that all K_p values showed fluctuations in air recovery. At K_p = 1.5, the system responded quickly but exhibited large variability, reflecting its increased sensitivity to noise and disturbances. In contrast, K_p = 0.5 resulted in a slower response, but fluctuations were still present. This suggests that lower proportional gains reduce the sensitivity of the control response but do not reduce the effect of noise. K_p = 0.8 provided a more balanced response, achieving reasonable response speed without the excessive variability observed at higher gains.

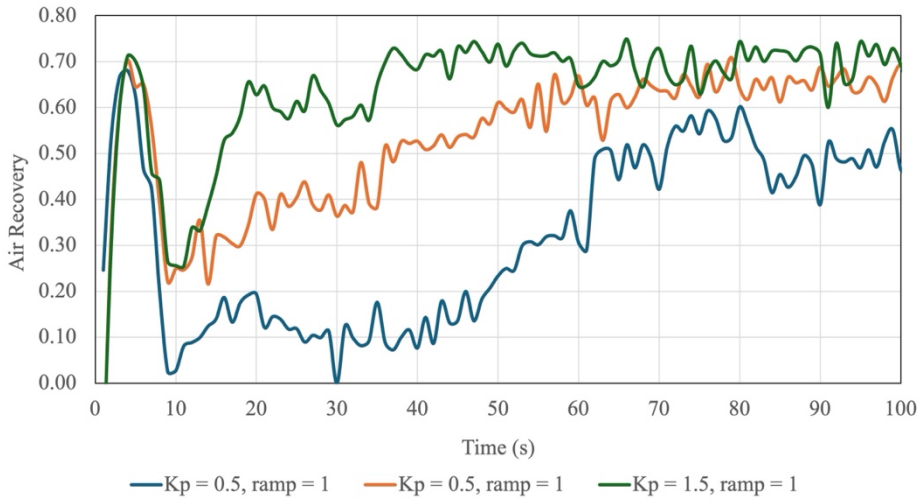


Figure 32: Air recovery response with ramp initialisation.

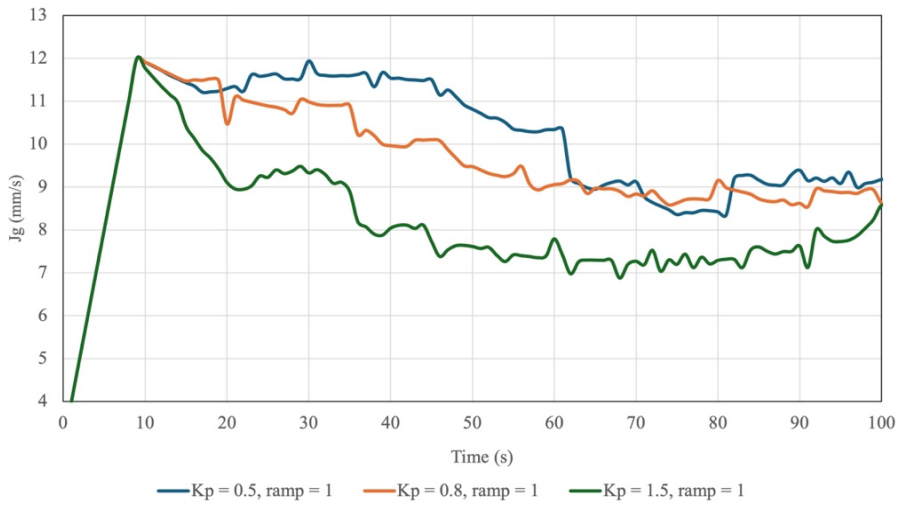


Figure 33: System response of superficial gas velocity with ramp initialisation.

As mentioned above, in previous tests, J_g was ramped at 0.5 mm/s, and the controller was activated after 10 steps at approximately 5 mm/s. Under these conditions, the air recovery converged to its target value more rapidly, reaching PAR in a few iterations.

In contrast, the current test used a faster ramp rate of 1 mm/s per step, which caused the controller to activate at a higher J_g of 12 mm/s. This delayed controller activation and resulted in the system progressing beyond the peak of the air recovery curve before the controller was activated. This resulted in a longer period of divergence before convergence to PAR. When the controller was activated, air recovery was substantially lower than in the slower ramp case, and the controller required more iterations to re-establish PAR. Across all gain values, the system converged more slowly and exhibited greater variability than under the slower ramp test cases.

These results highlight the significance of initial activation conditions. Controller activation closer to the PAR region reduced the number of iterations required for convergence and improved overall stability, while delayed activation at hight J_g values increased transient deviation and slowed the establishment of optimal conditions.

Figure 33 illustrates the J_g response under the same conditions. Similar trends were observed, where K_p = 1.5 led to a rapid response, rapidly reducing J_g after the ramp phase, with notable fluctuations. The lower K_p values resulted in a more gradual adjustment of J_g, with K_p = 0.5 displaying the slowest response. K_p = 0.8 demonstrated an intermediate behaviour, balancing responsiveness and control stability.

The system's response to an alternating step input pattern was analysed to assess its ability to reach and maintain PAR. The parameters for these tests are summarised in Table 11, while the results are illustrated in Figures 34 to 37. As outlined in Section 3.2.1, the regression model used the ten most recent data points to estimate the coefficients A, B, and C, discarding the oldest data points with each new iteration. This regression window remained active throughout the entire sequence, ensuring that coefficient estimation was always based on the latest process information.

The step sequence pattern involved a repeating series of increasing and decreasing J_g values. The increasing sequence progressed in a fixed increment of 4 mm/s (4, 8, 12 mm/s), while the decreasing sequence followed a variable pattern (9, 5, 3 mm/s) before repeating. This pattern continued until the tenth iteration, when the controller was activated. This step sequence enabled the evaluation of the controller's ability to converge on PAR under fluctuations in J_g, replicating fluctuating aeration behaviour typical in industrial flotation systems.

Table 11: Simulation parameters when superficial gas velocity is initialised through a repeating step sequence.

Parameter	Value	Unit
Froth height (h_f)	106	mm
Proportional controller gain (K_p)	0.5; 0.8; 1	-
Superficial gas velocity (J_g) sequences	4, 8, 12	mm/s
	9, 5, 3	mm/s

Figures 34 and 35 illustrate the response of α and J_g when subjected to an ascending J_g step sequence. Across all K_p values, the system approached and, in some cases, reached the PAR region after a few iterations. The differences in response characteristics varied depending on K_p. Higher gains produced faster convergence but also introduced larger fluctuations in both α and J_g. While lower gains resulted in a slower, more gradual convergence with reduced variability in the system's behaviour.

The slight negative α values observed at the start of the simulation arise from the unconstrained regression model and do not represent a meaningful result. This behaviour was retained to

preserve the model's full response range to evaluate the controller's behaviour under unconstrained conditions.

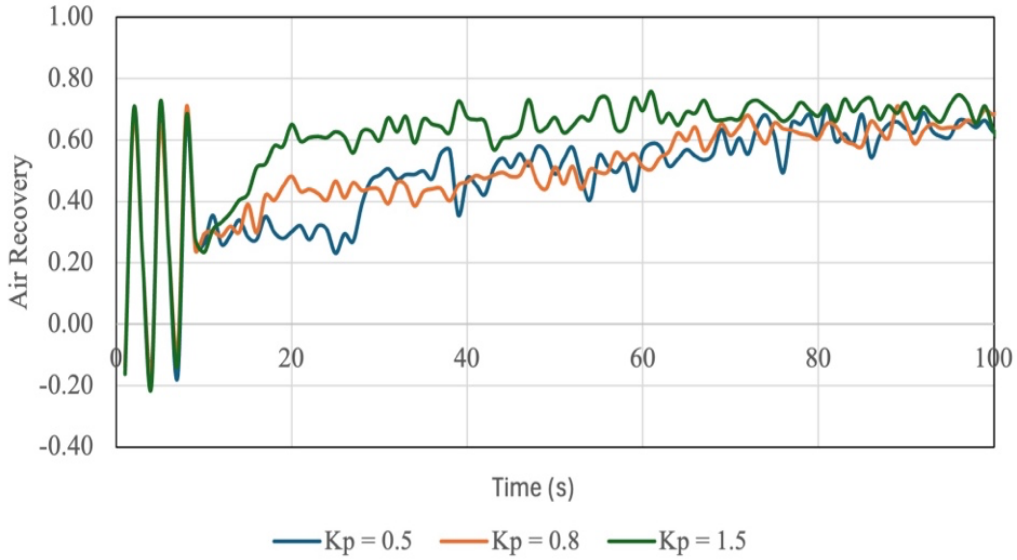


Figure 34: Air recovery trends in response to superficial gas velocity step input.

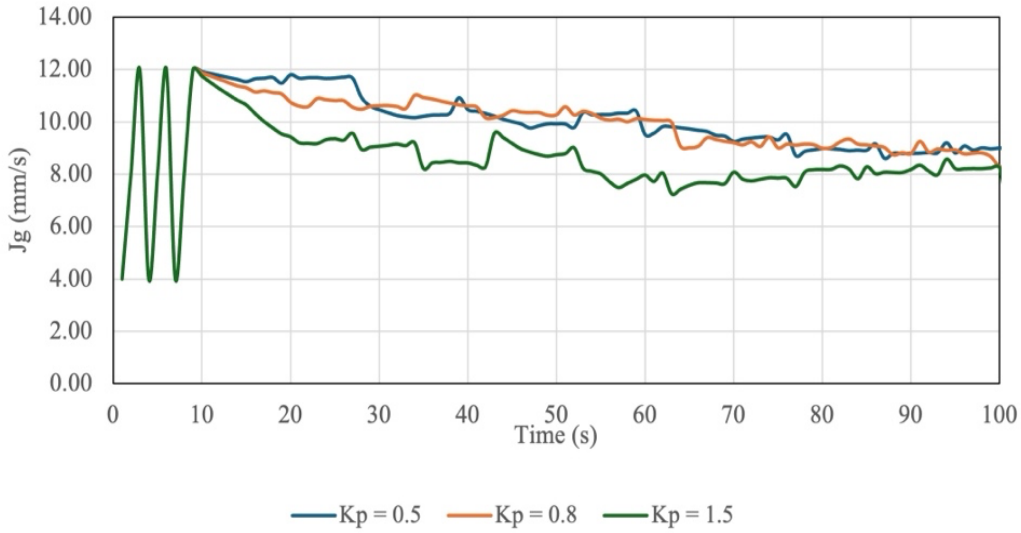


Figure 35: Superficial gas velocity response to a repeating step initialisation.

Figures 36 and 37 show the system's response under a decreasing J_g step sequence, where J_g was decreased from a value above its target value. This test assessed the controller's ability to converge on the optimal J_g when starting from an elevated initial value. In this test, the system reached PAR more effectively, achieving convergence within fewer iterations across all K_p

values. For K_p = 0.5 and 0.8, the convergence was controlled and consistent, while K_p = 1.5 showed more pronounced fluctuations but still achieved PAR.

The difference in final J_g values before controller activation further supports these trends. The ascending sequence ended at J_g = 12, which was significantly higher than the optimal J_g, whereas the descending sequence ended at J_g = 4, positioned much closer to the target J_g. This confirms that when the system begins closer to its optimal operating point, convergence to PAR occurs more rapidly and with less fluctuations, regardless of the gain value.

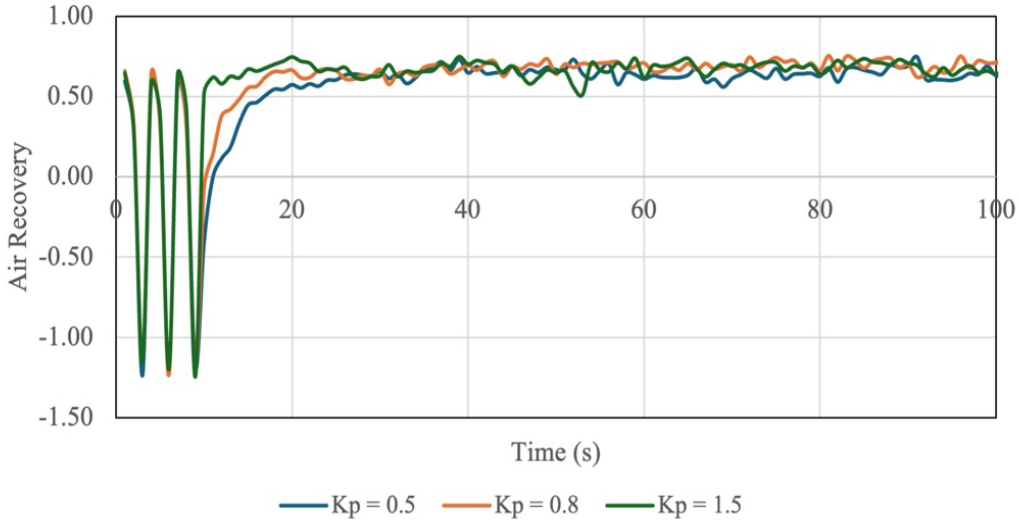


Figure 36: Air recovery response to decreasing superficial gas velocity step sequence.

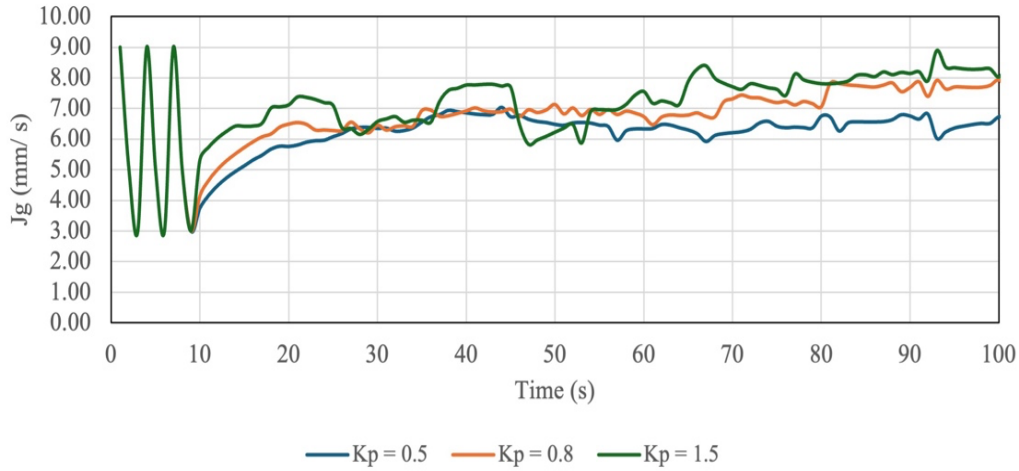


Figure 37: Superficial gas velocity response to a decreasing repeating step initialisation.

As previously mentioned, the tests discussed in this section were aimed at assessing how initial J_g values influenced the system's ability to converge on PAR. The repeating step sequence

initialisation proved to be more challenging than the ramp input, particularly when using an ascending step input (4 – 8 – 12 mm/s). Initialising J_g far from the optimal range (5 – 9 mm/s) resulted in larger corrective actions by the controller, which introduced instability in both α and J_g. As discussed, higher gains amplified this variability. In contrast, the descending step sequence (9 – 5 – 3 mm/s) brought the system closer to the target J_g before controller activation. When J_g was approximately at 5 mm/s, at the point of controller activation, the required corrections were smaller, and convergence to PAR occurred within fewer iterations.

When J_g started far from the desired range, such as J_g = 12 for the increasing step input, the controller faced significant challenges in dampening fluctuations and achieving steady state. The large difference between the initial and target J_g values led to larger corrective actions, which in turn caused greater fluctuations in both J_g and air recovery. Figures 34 and 35 show pronounced fluctuations immediately after the controller was activated, with these fluctuations being more significant at higher K_p values. Under this gain setting, the system took longer to stabilise and required more time for J_g to reach the optimal value.

These results highlight that the initial J_g plays an important role in determining convergence time and overall control performance. When the controller was activated near the optimal range, convergence was faster and more consistent, especially under moderate gain values. However, activation from a point far beyond PAR delayed convergence and introduced greater variability, highlighting the importance of selecting appropriate initialisation conditions to improve overall control performance.

4.3.6 System Response with Exponential Smoothing Filters

As discussed in Section 4.3., the regression based model occasionally exhibited ill-conditioning during proportional control simulations, resulting in unstable or inconsistent coefficient values. Although reverting to previously stable parameter values improved short term stability, this approach proved insufficient when numerical instability became more sever. To address this, exponential smoothing filters were introduced into the control framework to stabilise the air recovery values used in the feedback calculation.

The smoothing process reduced high frequency fluctuations in the measured signals, leading to a reduction in the propagation of noise and improving matrix conditioning during regression updates. In addition to smoothing, Tikhonov regularisation was also applied to reduce variability in the estimated coefficients and improve numerical stability. The combined effect of these techniques was evaluated through a number of simulation tests.

A series of simulations were conducted to evaluate the performance of different gain values and exponential smoothing filter gain configurations on the control system. These exponential smoothing filters, defined by Equation 18, applied weighted averaging through gain parameters G_1 and G_2, which determined the balance between current and historical data. These filters are widely used to reduce high frequency noise in control systems, the filter weightings tested are summarised in Table 12 which also summarises all the parameters used for the various tests. The outcomes of these simulations are shown in Figures 38 to 41.

Table 12: Parameters used for simulation with varying filter gain configurations.

Parameter	Value	Unit
Froth height (h_f)	106	mm
Noise standard deviation (σ^2)	0.001	
Proportional controller gain (K_p)	0.5; 0.8; 1.5	
Filter configurations		
Gain (G_1)	0.2; 0.4; 0.5; 0.6; 0.8	
Gain (G_2)	0.8; 0.6; 0.5; 0.4; 0.2	

Figure 38 shows the system's performance under a moderate noise level ($\sigma^2 = 0.001$) at different K_p values for the filter configuration $G_1 = 0.6$ and $G_2 = 0.4$. The graph shows that higher gains ($K_p = 1.5$) produced faster initial convergence toward the targeted air recovery value but also introduced significant fluctuations before stabilising near PAR. Conversely, lower gains ($K_p = 0.5$) demonstrated smoother but slower convergence to PAR.

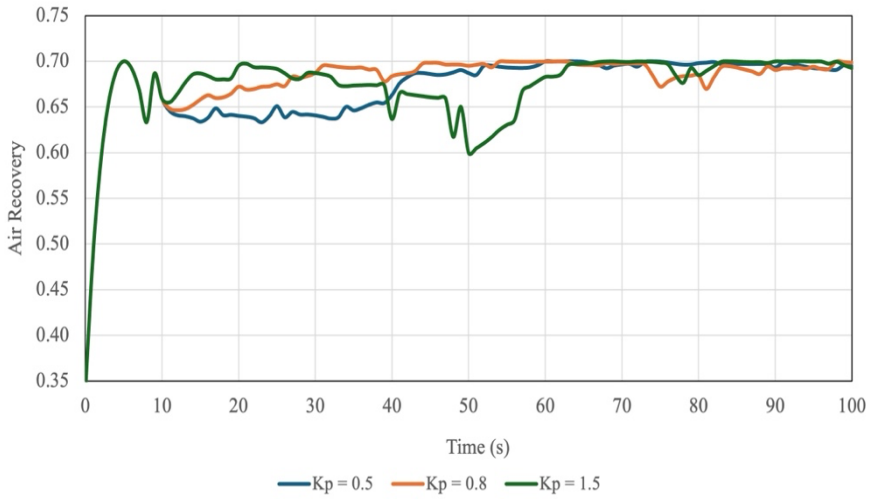


Figure 38: Air Recovery with filter weights of $G_1 = 0.6$, $G_2 = 0.4$

Figures 39 to 41 present the system responsiveness in terms of air recovery under these varying filter configurations across different K_p values at a moderate level of noise. Performance metrics for $K_p = 0.8$ are presented in Tables 12 and 13. Additional performance metrics for $K_p = 0.5$ and $K_p = 1.5$ are summarised in the Appendix in Tables A1 to A4, providing detailed insight into the behaviour of both air recovery and superficial gas velocity.

Table 13: Performance Metrics for air recovery at $K_p = 0.8$.

$K_p = 0.8$	Filter configuration				
Metric	$G_1 = 02$, $G_2 = 0.8$	$G_1 = 0.4$, $G_2 = 0.6$	$G_1, G_2 = 0.5$	$G_1 = 0.6$, $G_2 = 0.4$	$G_1 = 0.8$, $G_2 = 0.2$
MAE	0.01	0.02	0.00	0.01	0.02
SD	0.02	0.02	0.01	0.01	0.02
RMSE	0.01	0.01	0.00	0.01	0.02
IAE	1.35	1.73	0.45	0.70	1.99
CPI	52.44	40.97	157.92	100.80	35.56

Table 14: Performance Metrics for superficial gas velocity at $K_p = 0.8$.

$K_p = 0.8$	Filter configuration				
Metric	$G_1 = 02$, $G_2 = 0.8$	$G_1 = 0.4$, $G_2 = 0.6$	$G_1, G_2 = 0.5$	$G_1 = 0.6$, $G_2 = 0.4$	$G_1 = 0.8$, $G_2 = 0.2$
MAE	0.45	0.44	0.17	0.27	0.58
SD	0.61	0.39	0.26	0.33	0.60
RMSE	0.13	0.39	0.13	0.11	0.27
IAE	45.48	44.64	17.67	26.78	58.27
CPI	16.66	16.97	42.86	28.29	13.00

Figure 39 illustrates the simulation results at $K_p = 0.5$ for various filter configurations. The results indicate that filter Configurations where $G_1 < G_2$ achieved faster convergence to the steady state operating point. These filters prioritised recent data while still retaining some historical data, allowing the controller to respond effectively to process variations without amplifying noise. As summarised in Tables A2 and A3, the configuration $G_1 = 0.4$, $G_2 = 0.6$ provided the best overall performance, producing low error metrics (MAE, SD, RMSE = 0.01) and high CPI values, such as 93.29 for α and 26.18 for J_g. These results confirm that moderate smoothing offers an optimal balance between responsiveness and stability.

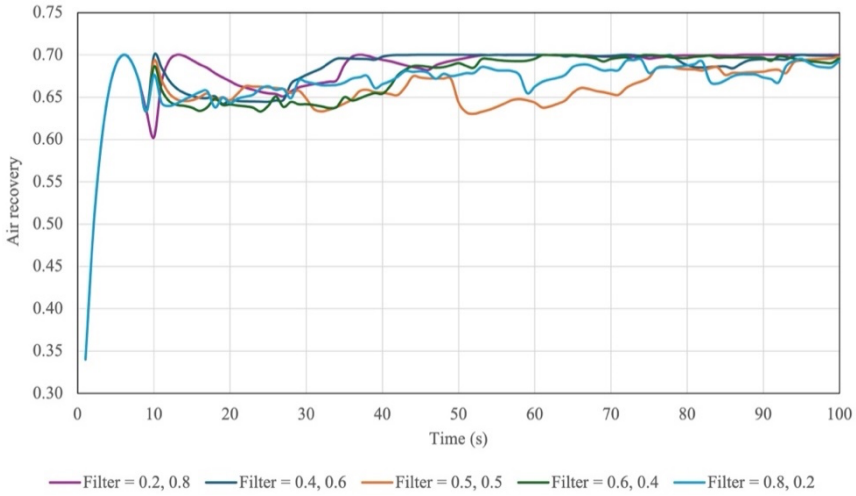


Figure 39: Air recovery response with $K_p = 0.5$ for varying filter gain values.

In contrast, configurations where $G_1 > G_2$ exhibited slower convergence and reduced responsiveness. The configuration $G_1 = 0.8$, $G_2 = 0.2$ produced higher error and variability, where the IAE for α increased to 1.50, accompanied by a decrease in CPI, 47.26 for α and

19.23 for J_g. This indicates that increased reliance on historical data overly dampens controller action, resulting in a delayed adjustment to changes in the process.

Intermediate configurations such as $G_1 = G_2 = 0.5$ also showed acceptable performance but were slightly less efficient, with lower CPI values, 27.76 for α and 13.56 for J_g. Overall, the results demonstrated that for $K_p = 0.5$, filters where $G_1 < G_2$ enhanced convergence speed and control accuracy, validating their suitability for mitigating noise while also preserving responsiveness under this gain value.

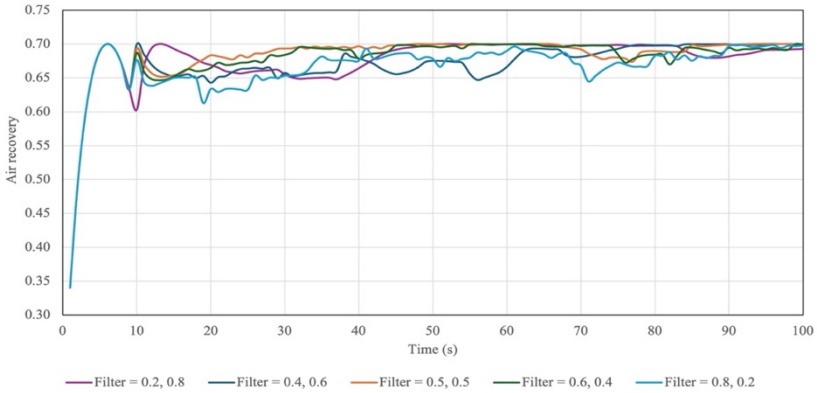


Figure 40: Air recovery response with $K_p = 0.8$ for different filter gain values.

Figure 40 presents the α results for $K_p = 0.8$ for various filter configurations. Configurations where $G_1 = G_2 = 0.5$ and $G_1 = 0.6$, $G_2 = 0.4$ were the most effective in achieving rapid and stable convergence toward the PAR value. These configurations provided a balance between responsiveness and noise suppression by weighting recent and historical data almost equally. As summarised in Tables 13 and 14, these configurations had the highest CPI values, 157.92 for α and 42.86 for J_g, and the lowest IAE and MAE values, indicating fast and accurate tracking of the target air recovery point with minimal fluctuations.

Configurations such as $G_1 = 0.2$, $G_2 = 0.8$, and $G_1 = 0.8$, $G_2 = 0.2$, resulted in slower or more variable convergence, reflected in their low CPI (52.44 and 35.56 for α; 16.66 and 13.00 for J_g) and high IAE values, where the high IAE values signify reduced control efficiency and increased settling behaviour. Although MAE and RMSE values remained low across all configurations for air recovery, SD values were the lowest for $G_1 = G_2 = 0.5$ and $G_1 = 0.6$, $G_2 = 0.4$ configurations, confirming that these settings offered the most consistent control response.

Overall, the results demonstrate that for higher gains, moderate smoothing, where filter gains are equal or slightly weighted toward recent data, optimises control performance by maintaining stability without sacrificing convergence speed.

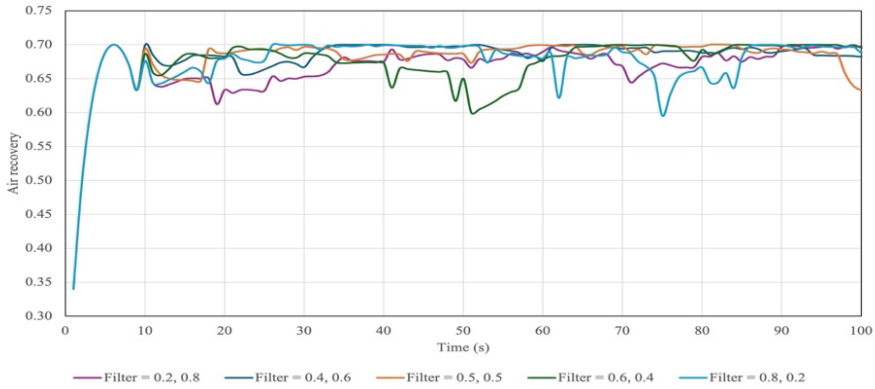


Figure 41: Air recovery response for various filter gain values for Kp = 1.5.

Figure 41 illustrates the air recovery response for various filter configurations at $K_p = 1.5$. All filter configurations exhibited increased variability in air recovery, consistent with the system's increased sensitivity at higher gain values. The configurations, $G_1 = 0.4$, $G_2 = 0.6$ performed the best, achieving convergence toward the PAR value. This configuration demonstrated the most favourable performance across all metrics producing low error values and high CPI values, 77.64 for α and 20.78 for J_g, indicating accurate and efficient control despite the inherent instability associated with higher gains.

The metrics summarised in Tables 13 and 14 as well as those in the Appendix, Tables A1 – A4, align well with the graphical trends. Filter configurations where $G_1 < G_2$, which placed higher weighting on recent data, enhanced responsiveness and enabled faster convergence but also introduced increased variability, especially at higher proportional gain values such as $K_p = 1.5$. In contrast, configurations where $G_1 > G_2$ placed more weighting on historical data, effectively dampening fluctuations but at the cost of slower convergence. This trade off proved useful under low gains, such as $K_p = 0.5$ where stability was prioritised over responsiveness.

While performance varied across different proportional gain values, at $K_p = 0.8$, configurations with nearly equal weighting, such as $G_1 = G_2 = 0.5$ and, $G_1 = 0.6$, $G_2 = 0.4$, offered the best compromise between noise suppression and convergence behaviour. These configurations achieved faster convergence to PAR with minimal variability in both air recovery and superficial gas velocity. This was reflected in the high CPI values and low cumulative errors

observed under these simulation conditions, confirming efficient control performance under moderate noise conditions

However, at other gains, different filter configurations performed better, indicating that filter effectiveness is closely linked to the selected gain value and should be tuned accordingly. Filters with $G_1 < G_2$ performed well under $K_p = 0.5$, offering reliable control with slightly slower convergence, while $K_p = 1.5$ was more sensitive to filter selection, with greater variability observed across all configurations due to amplified gin induced sensitivity.

Across all test conditions, $K_p = 0.8$ returned the highest CPI values, indicating superior control performance in terms of convergence speed, tracking accuracy, and robustness to noise. These findings suggest that $K_p = 0.8$ offers the best trade-off between responsiveness and control variability and represents the optimal proportional gain for maintaining a stable yet responsive control loop.

Based on the metrics and graphs, the best performing combination for achieving PAR and maintaining control performance was $K_p = 0.8$ with the filter configuration $G_1 = 0.6$, $G_2 = 0.4$. This configuration produced the highest CPI, moderate MAE and SD, and fast convergence to steady state, indicating strong control efficiency. For processes prioritising low variability over responsiveness, $K_p = 0.5$ with a configuration of $G_1 = 0.4$, $G_2 = 0.6$ also performed well, offering low error and sufficient dampening of fluctuations.

Overall, the implementation of exponential smoothing filters and regularisation noticeably improved controller stability and responsiveness across all tested gain values. The analysis demonstrated that appropriate filter tuning can mitigate the effects of measurement noise and numerical sensitivity, especially in systems susceptible to ill-conditioning. The results also indicate that the optimal filter configuration depends on the selected proportional gain, emphasising the importance of tuning between the controller and filter parameters. These results provide a strong foundation for developing more advanced gain scheduling approaches that dynamically adjust filter weightings to maintain consistent performance under varying process and disturbance conditions.

4.4 Cost Function Analysis and Controller Tuning

Following the methodology described in Section 3.3.5 the proportional gain was optimised using the cost function defined in Equation 20, which quantifies controller performance as the sum of square errors between the system output (air recovery) and the target peak air recovery value. Two approaches were implemented:

- Parameter Sweep: K_p was varied between 0.20 and 1.50 in increments of 0.05, with each simulation using the same noise level to ensure consistency. The SSE and SD of air recovery at steady state were recorded.
- Iterative tuning: starting from an initial K_p, $K_p = 0.20$, the value was adjusted in steps, reversing and halving the step size whenever performance did not improve. The K_p adjustment continued until convergence was achieved according to the stopping criteria.

The sweep results in Table A6 showed that $K_p = 0.85$ yielded the lowest SSE (2.73×10^{-4}) and a low SD (1.12×10^{-3}), indicating a good balance between control accuracy and stability. Several gains, $K_p = 0.65$, 0.70, and 0.75, also performed well, having slightly higher SSE but similar SD values.

The iterative tuning method converged to $K_p = 0.70$ with an SSE of 2.63×10^{-4} and an SD of 2.49×10^{-3}. While this SSE was slightly lower than the sweep method, the difference was not statistically significant. The sweep method evaluates all gains within the specified range, whereas the iterative tuning method adjusts its direction based on improvements and converges to a different point within the region of near optimal performance.

Both methods confirm that the optimal K_p lies between 0.70 and 0.85 under the tested filter configuration ($G_1 = 0.6$, $G_2 = 0.4$) and noise conditions. The similar SSE values indicate that either approach is suitable for this system, but the sweep method provided a more comprehensive view of the relationship between K_p and controller performance.

4.5 Disturbance Rejection Performance of the P Controller

The disturbance rejection tests presented in this section were conducted using the cost function defined in Section 4.4. Although the control model operates in a steady state framework,

disturbances were introduced in the input variables to simulate realistic process deviations and assess how effectively the proportional controller establishes PAR. It was found that intermediate K_p values achieved the best trade-off between responsiveness, dampening of fluctuations, and control effort. These findings confirm that this tuning range provides robust performance for disturbance rejection within steady state control conditions.

In all four disturbance scenarios, the gains producing the lowest SSE values fell within this same range, reinforcing that the tuning methodology developed earlier is transferable to real time control simulations involving process disturbances. Table 15 summarises the results from the simulations conducted with Figures 42 to 46 illustrating the system responses. As outlined in Section 3.3.3.7, each disturbance rejection simulation was repeated twenty times using randomly generated noise. The K_p, SSE, and SD values presented in Table 15 represent the averages calculated from repeated runs, providing statistically reliable measures of controller performance under noisy simulation conditions

The tests aim to evaluate the controller's ability to maintain flotation performance when subjected to significant process disturbances, using the following four scenarios:

1. An abrupt decrease in superficial gas velocity to 5 mm/s.
2. An abrupt increase in froth height from 106 to 112 mm.
3. A gradual decrease in froth height from 106 to 98 mm.
4. A gradual increase in froth height from 106 112 mm.

The controller performance was quantified using SSE, SD and settling time. The best performing K_p value for each case was determined based on the lowest SSE.

Table 15: Disturbance rejection test results

Disturbance type	Best Kp	SSE	SD	Settling time (s)
Abrupt J_g drop to 5 mm/s	0.70	3.52x10-4	0.0264	90
Abrupt h_f increase to 112 mm	0.80	3.37x10-4	0.00117	200
Gradual h_f decrease to 98 mm	0.90	3.64x10-4	0.0233	229
Gradual h_f increase to 112 mm	0.70	3.56x10-4	0.0132	210

The settling times reported in Table 15 were obtained automatically using MATLAB's stepinfo function, which identifies the point at which a variable remains within a $\pm$ 2% band of its final steady state value. Much like the steady state simulations previously conducted, the time axis represents iteration steps rather than physical time. Therefore, the reported values reflect the number of iterations required for the response to remain within this tolerance. Although the air recovery responses appear visually stable as illustrated in Figures 44 and 45, small residual fluctuations and numerical noise caused by MATLAB to register longer settling durations. The reported settling times therefore provide a quantitative measure of damping performance rather than a physical process time.

4.5.1 Abrupt Superficial Gas Velocity Reduction System Response

Figure 42 shows the system's response to a sudden drop in superficial gas velocity to 5 mm/s at t = 200 s. The disturbance caused an immediate decrease in air recovery, but the controller restored both the superficial gas velocity and air recovery to their previous levels before the disturbance within approximately 90 iterations. At the optimal tuning ($K_p = 0.7$), the simulation achieved an SSE of 3.52×10^{-4}, and an SD of 0.0264, confirming that the proportional controller effectively mitigated abrupt process disturbances.

This behaviour mirrors the results from Section 4.3.2, where a moderate K_p value ($K_p = 0.8$), achieved the best performance under ideal conditions, delivering fast recovery without fluctuations or variability. The relatively short settling time compared to the froth height disturbance demonstrates the controller's ability to reject disturbances in the airflow effectively, this being confirmed by the moderate SD value. The SD value indicates that while the recovery was quick, there was some variability during convergence, which is acceptable because of the magnitude of the disturbance.

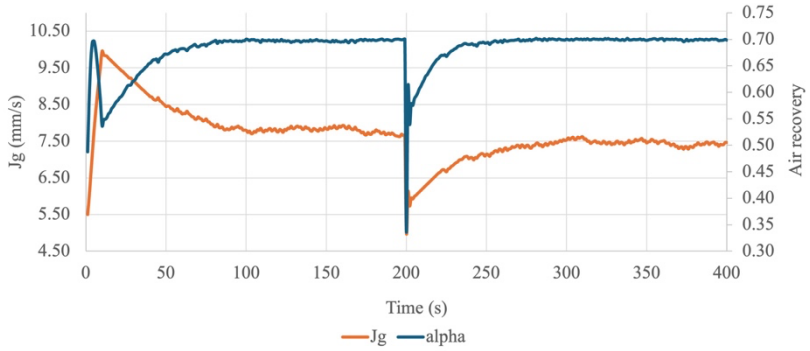


Figure 42: System response to an abrupt decrease in the superficial gas velocity at t = 200 s.

4.5.2 Abrupt Froth Height Increase System Response

Figure 43 illustrates the effect of a sudden increase in froth height from 106 to 112 mm. The disturbance resulted in a sharp decrease in air recovery stabilising at a new steady state of 0.58. This behaviour is consistent with the trend observed in Figure 28, confirming that the change in froth height shifted the position of PAR. The superficial gas velocity remained unchanged, as the controller maintained the system at the new PAR corresponding to the altered froth height.

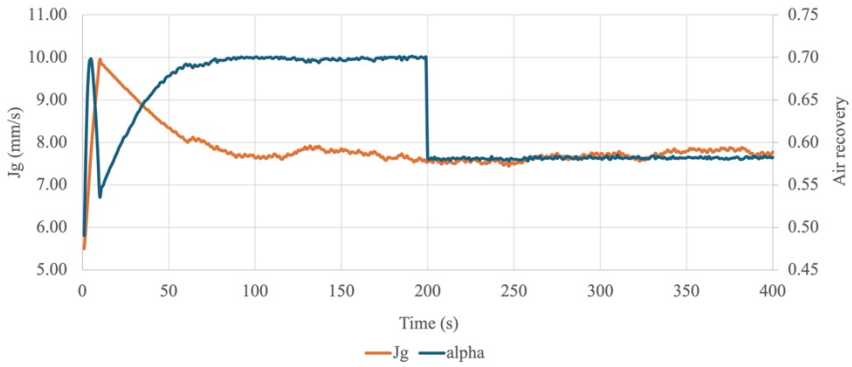


Figure 43: System response to an abrupt increase in froth height from 106 to 112 mm at t = 200 s.

This behaviour is consistent with the steady state results in Section 4.3.2, which showed that increasing froth height increases bubble residence time and the probability of coalescence or bursting before reaching the froth surface. Although the simulation in this section operates in an iterative steady state framework, the same process relationships govern the response The change in froth height caused a slight shift in the position of PAR. This difference falls within the range of variability introduced by measurement noise, indicating that the controller maintained the system near the new PAR rather than restoring the previous one. A K_p of 0.8 was calculated by the system to be the best gain, which achieved the lowest SSE (3.37×10^{-4})

and a low SD (0.00117), reflecting the controller's ability to remain consistent control performance once the new operating point was reached.

4.5.3 Gradual Froth Height Increase System Response

Figure 45 shows the system's response to a gradual increase in froth height from 106 to 112 mm in increments of 0.5 mm. The disturbance produced a steady decline in air recovery until it reached a new steady state (0.58). However, superficial gas velocity was not affected by this change. As with the abrupt increase in froth height simulation, the system stabilised at a lower PAR and the controller could not counteract the effect of the increased froth height through superficial gas velocity adjustments alone.

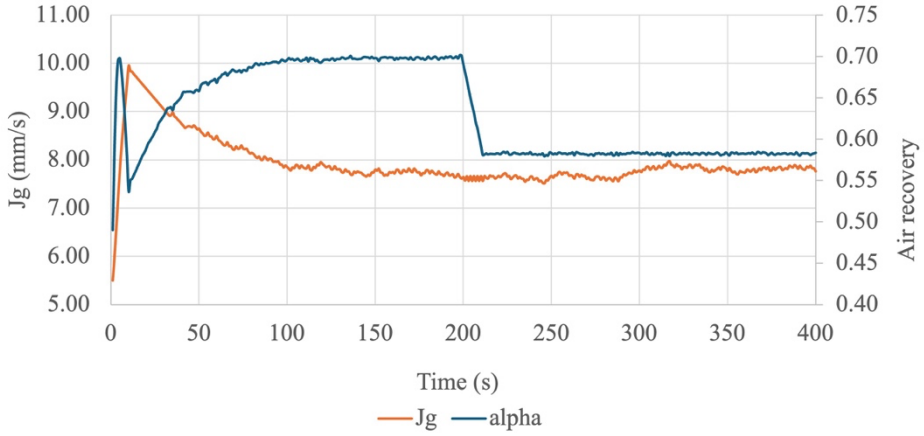


Figure 44: System response to a gradual increase in froth height from 106 to 112 mm at t = 200 s.

The best performing gain (K_p = 0.70), remained within the optimal range identified in Section 4.4 and produced an SSE of 3.56 x 10^{-4} and a low SD of 0.0132, indicating a consistent response with minimal variability. Similar to the abrupt froth height increase test, the relatively long settling time (210 iteration steps), reflects the gradual nature of the applied disturbance, with the final steady state only being reached after the last increment in froth height.

4.5.4 Gradual Froth Height Decrease System Response

A simulation was conducted where froth height was gradually decreased by increments of 0.5 mm from 106 to 98mm. Figure 44 shows that air recovery increased progressively during the reduction in froth height, while the controller adjusted J_g, where the change in J_g resulted from the controller's feedback action. According to the process relationship, a decrease in froth

height reduces the froth residence time, which in turn shifts the PAR to a higher value. The controller therefore responded to the change by adjusting J_g in order for the process to reach the new PAR point.

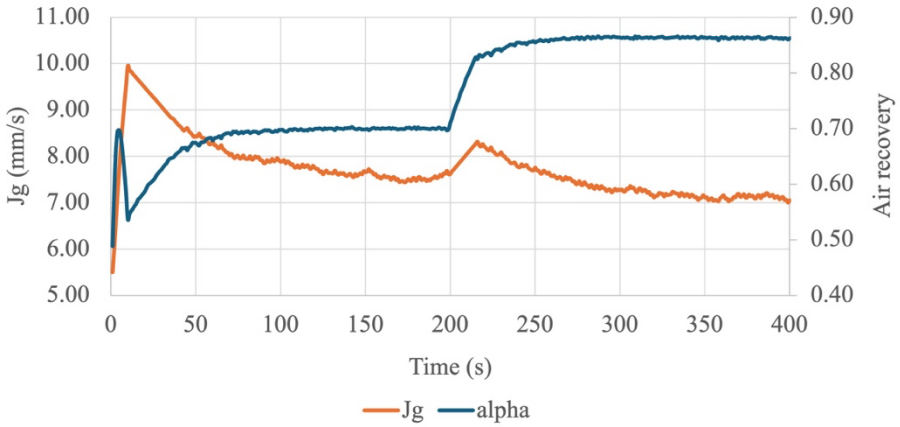


Figure 45: System response to a gradual decrease in froth height at t = 200 s from 106 to 98 mm.

The best performance was achieved by a gain of $K_p = 0.9$, which produced the lowest SSE while maintaining acceptable variability. The slightly higher SD compared with the froth height increase test occurred because the controller continued to respond to small fluctuations during the gradual disturbance, introducing minor variations as the system approached its new steady state region. These results show that moderate proportional gains enable effective adaptation to process induced shifts in the PAR point.

The disturbance rejection results, along with the steady state and noise performance analyses in earlier sections, provide a complete picture of the P controller's capabilities and limitations in maintaining a system at PAR. The findings confirm that the tuning methodology developed using the cost function is robust across a range of operating conditions, while also highlighting process scenarios where the controller was insufficient. Chapter 5 summaries the study's key findings and discusses the limitations of the current approach including potential for future work.

CHAPTER 5

5 Conclusion

This chapter brings together the key findings of the study, discussing the overall results from the steady state analysis, noise sensitivity testing, and disturbance rejection performance to provide a full view of the proportional control strategy used for maintaining flotation cells at the point of Peak Air Recovery. This chapter also examines the limitations observed, especially in handling froth height disturbances that shift the process optimum point. It also considers how these limitations could be addressed through enhanced control methods.

The aim of this research was to develop and evaluate a control strategy that could effectively maintain cells at PAR by focusing on air recovery optimisation, with the effect on valuable mineral recovery inferred through its established correlation with PAR and froth stability. This began with the development of a model for air recovery as a function of froth height and superficial gas velocity under both steady state and dynamic conditions. The model, based on regression derived parameters from parameters from literature, was supplemented with an analytical method. This analytical work offered a theoretical benchmark for validating simulation behaviour in noise free conditions.

Building on this modelling framework, a robust proportional control strategy was evaluated using the cost function described in Section 4.4. The cost function performed a grid search sweep over predefined proportional gain values and selected the gain that minimised the cost function under identical noise and filter settings. This optimisation method consistently yielded gains in the 0.70 and 0.90 range, with $K_p = 0.85$ giving the lowest average cost under moderate noise conditions reaching the target PAR in approximately 50 iterations with minimal overshoot and low variability.

Filter configuration also had a significant influence on the system behaviour, affecting responsiveness and noise suppression. The results in Section 4.4 identified $G_1 = G_2 = 0.50$ as the highest ranked configuration, offering the best overall balance across performance metrics

at K_p = 0.80. However, G_1 = 06, G_2 = 0.4 was selected for the remainder of the simulations. This configuration achieved CPI values close to the best ranked configuration, while also enabling fast adaptation to process shifts by placing more weight on recent process measurements ($G_1 > G_2$), thereby supporting timely corrective action without excessive noise sensitivity. Equal weighting ($G_1 = G_2$) was less responsive in some cases, as older data continued to have equal influence. This potentially delayed correction, while configurations placing more weight on older data risked amplifying high frequency noise. In this way, the selected configuration balanced the trade-off between responsiveness and noise suppression, improving the controller's ability to track changes while maintain stability.

The performance of this strategy was further quantified through convergence time, SSE, SD, overshoot, and stability, including tests under realistic process disturbances. Abrupt reductions in superficial gas velocity were rejected quickly, restoring air recovery within 90 iterations to PAR levels before disturbances were introduced. Gradual decreases in froth height led to a steady increase in air recovery, with the controller temporarily increasing the superficial gas velocity before settling at a new steady state. In both abrupt and gradual froth height increases, air recovery settled at a lower value and superficial gas velocity remained unchanged, confirming that such changes shift the process optimum and cannot be compensated by manipulating the superficial gas velocity alone. Across all disturbance test cases, the best performing K_p values fell within the same 0.70 to 0.90 range identified earlier, and average results from 20 simulations per test case closely matched the cost function optimum of K_p = 0.85.

When compared with alternative strategies, P controller demonstrated distinct advantages. ESC, while adaptive, converges slowly due to its perturbation based search, a limitation that is amplified in industrial settings where metallurgical sampling intervals for feed, concentrate, and tailings are often 20 to 60 minutes to obtain representative data. Under such conditions, ESC may determine a new optimum only after feed conditions have already been changed. In contrast, P controller converged in under a minute in numerous test cases, ensuring that adjustments were computed well within practical feedback cycles. MPC offered predictive capabilities but required significantly greater computational resources.

The findings highlight that a proportional controller, when tuned using a cost function-based approach with an appropriate exponential smoothing filter, can maintain PAR effectively

across a wide range of operating conditions, however, several limitations remain. The fixed tuning approach means that long term process drift or significant changes in ore characteristics may require retuning, froth disturbances that shift the process optimum cannot be fully compensated without directly manipulating froth height, and while the chosen filter balanced responsiveness and noise suppression well, extreme process noise levels or faulty measurements could still cause delayed correction or instability.

Future development should focus on adaptive tuning of the filter configuration to maintain performance under changing conditions, integrating froth height control with superficial gas velocity control to address shifts in the process peak, and validating the strategy at plant scale under variable feed and reagent conditions.

In conclusion, the study achieved its aim and addressed its objectives, demonstrating that a well-tuned proportional controller can deliver fast, consistent, and practical control of PAR. Its ability to operate within standard industry sampling intervals, positions it as a viable option for real time flotation control. This work not only demonstrates that a simple, well-tuned proportional controller can deliver fast and stable PAR regulation but also establishes a foundation on which more advanced adaptive and multi variable control systems can be built for future industrial application.

REFERENCES

Abdul-Adheem, W., Riyadh Abdul-Adheem, W., 2017. From PID to Nonlinear State Error Feedback Controller, Article in International Journal of Advanced Computer Science and Applications.

Barbian, N., Hadler, K., Cilliers, J.J., 2006. The froth stability column: Measuring froth stability at an industrial scale. Miner Eng 19, 713–718. https://doi.org/10.1016/j.mineng.2005.09.021

Barbian, N., Hadler, K., Ventura-Medina, E., Cilliers, J.J., 2005. The froth stability column: linking froth stability and flotation performance. Miner Eng 18, 317–324. https://doi.org/10.1016/j.mineng.2004.06.010

Barbian, N., Ventura-Medina, E., Cilliers, J.J., 2003. Dynamic froth stability in froth flotation. Miner Eng 16, 1111–1116. https://doi.org/10.1016/j.mineng.2003.06.010

Barry A. Wills, James A. Finch, 2016. Wills' Mineral Processing Technology. pp. 265–380. https://doi.org/DOI:10.1016/B978-0-08-097053-0.00012-1

Barry A. Wills, James A. Finch, 2015. Wills' Mineral Processing Technology An Introduction to the Practical Aspects of Ore Treatment and Mineral Recovery Eighth Edition. https://doi.org/https://doi.org/10.1016/C2010-0-65478-2

Bergh, L.G., Yianatos, J.B., 2011. The long way toward multivariate predictive control of flotation processes. J Process Control 21, 226–234. https://doi.org/10.1016/J.JPROCONT.2010.11.001

Brooks, K., Munalula, W., 2017. Flotation Velocity and Grade Control Using Cascaded Model Predictive Controllers, in: IFAC-PapersOnLine. Elsevier B.V., pp. 25–30. https://doi.org/10.1016/j.ifacol.2017.12.005

Brooks, K., Westcott, M., Bauer, M., 2019. A Combined MPC for Milling and Flotation – A Simulation Study. IFAC-PapersOnLine 52, 24–29. https://doi.org/10.1016/j.ifacol.2019.09.158

Brooks, K.S., Ramodike, K., Higginson, A., 2024. Controlling a Flotation Cell at the Peak Air Recovery Point, in: IFAC-PapersOnLine. Elsevier B.V., pp. 13–17. https://doi.org/10.1016/j.ifacol.2024.09.283

Cai, X., Sun, P., Chen, J., Xie, L., 2014. Rapid distributed model predictive control design using singular value decomposition for linear systems. J Process Control 24, 1135–1148. https://doi.org/10.1016/j.jprocont.2014.05.002

Cui, C., Wang, H., Wei, Y., 2023. Perturbations of Moore-Penrose inverse and dual Moore-Penrose generalized inverse. J Appl Math Comput 69, 4163–4186. https://doi.org/10.1007/s12190-023-01920-5

Deglon, D.A., 2005. The effect of agitation on the flotation of platinum ores, in: Minerals Engineering. pp. 839–844. https://doi.org/10.1016/j.mineng.2005.01.024

Fuerstenau, M.C., Jameson, G.J., Yoon, R.-H., 2007. Froth Flotation: A Century of Innovation. SME.

Gharai, M., Venugopal, R., 2016. Modeling of Flotation Process - An Overview of Different Approaches. Mineral Processing and Extractive Metallurgy Review 37, 120–133. https://doi.org/https://doi.org/10.1080/08827508.2015.1115991

Hadler, K., Cilliers, J.J., 2009a. The relationship between the peak in air recovery and flotation bank performance. Miner Eng 22, 451–455. https://doi.org/10.1016/j.mineng.2008.12.004

Hadler, K., Cilliers, J.J., 2009b. The relationship between the peak in air recovery and flotation bank performance. Miner Eng 22, 451–455. https://doi.org/10.1016/j.mineng.2008.12.004
Hadler, K., Greyling, M., Plint, N., Cilliers, J.J., 2012a. The effect of froth depth on air recovery and flotation performance. Miner Eng 36–38, 248–253. https://doi.org/10.1016/j.mineng.2012.04.003
Hadler, K., Greyling, M., Plint, N., Cilliers, J.J., 2012b. The effect of froth depth on air recovery and flotation performance. Miner Eng 36–38, 248–253. https://doi.org/10.1016/J.MINENG.2012.04.003
Hadler, K., Smith, C.D., Cilliers, J.J., 2010a. Recovery vs. mass pull: The link to air recovery, in: Minerals Engineering. pp. 994–1002. https://doi.org/10.1016/j.mineng.2010.04.007
Hadler, K., Smith, C.D., Cilliers, J.J., 2010b. Recovery vs. mass pull: The link to air recovery. Miner Eng 23, 994–1002. https://doi.org/10.1016/J.MINENG.2010.04.007
Hodouin, D., 2011. Methods for automatic control, observation, and optimization in mineral processing plants. J Process Control 21, 211–225. https://doi.org/10.1016/j.jprocont.2010.10.016
Hoerl, A.E., Kennard, R.W., 1970. Ridge Regression: Biased Estimation for Nonorthogonal Problems. Technometrics 12, 55–67. https://doi.org/https://doi.org/10.2307/1267351
Igual, C., Igual, J., 2024. Simultaneous Three-Degrees-of-Freedom Prosthetic Control Based on Linear Regression and Closed-Loop Training Protocol. Sensors 24. https://doi.org/10.3390/s24103101
Jovanović, I., Miljanović, I., 2015. Contemporary advanced control techniques for flotation plants with mechanical flotation cells – A review. Miner Eng 70, 228–249. https://doi.org/10.1016/j.mineng.2014.09.022
Katsikis, V.N., Pappas, D., Petralias, A., 2011. An improved method for the computation of the Moore-Penrose inverse matrix. Appl Math Comput 217, 9828–9834. https://doi.org/10.1016/j.amc.2011.04.080
Kawatra, S., 2011. Fundamental principles of froth flotation. SME Mining Engineering Handbook 2, 1517–1531.
Köme, C., Yazlik, Y., 2024. Inverse and Moore–Penrose inverse of conditional matrices via convolution. J Appl Math Comput 70, 417–433. https://doi.org/10.1007/s12190-023-01974-5
Leiva, J., Vinnett, L., Yianatos, J., 2012. Estimation of air recovery by measuring froth transport over the lip in a bi-dimensional flotation cell. Miner Eng 36–38, 303–308. https://doi.org/10.1016/j.mineng.2012.05.023
Liu, J.J., MacGregor, J.F., 2008. Froth-based modeling and control of flotation processes. Miner Eng 21, 642–651. https://doi.org/10.1016/j.mineng.2007.12.011
Maldonado, M., Araya, R., Finch, J., 2012. An overview of optimizing strategies for flotation banks. Minerals 2, 258–271. https://doi.org/10.3390/min2040258
Mondal, S., Acharjee, A., Mandal, U., Saha, B., 2021. Froth flotation process and its application. Vietnam Journal of Chemistry. https://doi.org/10.1002/vjch.202100010
Nakhaei, F., Mosavi, M.R., Sam, A., Vaghei, Y., 2012. Recovery and grade accurate prediction of pilot plant flotation column concentrate: Neural network and statistical techniques. Int J Miner Process 110–111, 140–154. https://doi.org/10.1016/j.minpro.2012.03.003
Neethling, S.J., Cilliers, J.J., 2008. Predicting and correcting grade-recovery curves: Theoretical aspects. Int J Miner Process 89, 17–22. https://doi.org/10.1016/j.minpro.2008.08.001

Neethling, S.J., Cilliers, J.J., 2003. Modelling flotation froths. Int J Miner Process 72, 267–287. https://doi.org/10.1016/S0301-7516(03)00104-2
Norori-McCormac, A., Brito-Parada, P.R., Hadler, K., Cole, K., Cilliers, J.J., 2017. The effect of particle size distribution on froth stability in flotation. Sep Purif Technol 184, 240–247. https://doi.org/10.1016/j.seppur.2017.04.022
Oosthuizen, D.J., Craig, I.K., 2019. Predicting optimal operating points by modelling different flotation mechanisms. IFAC-PapersOnLine 52, 60–65. https://doi.org/10.1016/j.ifacol.2019.09.164
Oosthuizen, D.J., Le Roux, J.D., Craig, I.K., 2024. Non-linear Model Predictive Control to Improve the Mineralogical Efficiency of Flotation Circuits, in: IFAC-PapersOnLine. Elsevier B.V., pp. 1–6. https://doi.org/10.1016/j.ifacol.2024.09.281
Oosthuizen, D.J., le Roux, J.D., Craig, I.K., 2021. A dynamic flotation model to infer process characteristics from online measurements. Miner Eng 167. https://doi.org/10.1016/j.mineng.2021.106878
Oosthuizen, D.J., Williams, B.A., van der Spuy, D.D.V., 2023. Barriers to Integrated Flotation Control Solutions: Lessons from an Industrial Implementation, in: IFAC-PapersOnLine. Elsevier B.V., pp. 2335–2340. https://doi.org/10.1016/j.ifacol.2023.10.1203
Ostadrahimi, M., Farrokhpay, S., Gharibi, K., Dehghani, A., 2021. Effects of operating parameters on the froth and collection zone recovery in flotation: An industrial case study in a 10 m3 cell. Minerals 11. https://doi.org/10.3390/min11050494
Qu, X., Wang, L., Nguyen, A. V., 2013. Correlation of air recovery with froth stability and separation efficiency in coal flotation. Miner Eng 41, 25–30. https://doi.org/10.1016/j.mineng.2012.10.013
Quintanilla, P., Neethling, S.J., Brito-Parada, P.R., 2021. Modelling for froth flotation control: A review. Miner Eng 162, 106718. https://doi.org/10.1016/j.mineng.2020.106718
Rahman, R.M., Ata, S., Jameson, G.J., 2015. Study of froth behaviour in a controlled plant environment – Part 1: Effect of air flow rate and froth depth. Miner Eng 81, 152–160. https://doi.org/10.1016/j.mineng.2015.07.003
Rahman, R.M., Ata, S., Jameson, G.J., 2012. The effect of flotation variables on the recovery of different particle size fractions in the froth and the pulp. Int J Miner Process 106–109, 70–77. https://doi.org/10.1016/j.minpro.2012.03.001
Remes, A., Saloheimo, K., Jämsä-Jounela, S.L., 2007. Effect of speed and accuracy of on-line elemental analysis on flotation control performance. Miner Eng 20, 1055–1066. https://doi.org/10.1016/j.mineng.2007.01.016
Schwarz, S., Alexander, D., 2006. Gas dispersion measurements in industrial flotation cells, in: Minerals Engineering. pp. 554–560. https://doi.org/10.1016/j.mineng.2005.09.022
Shean, B., Hadler, K., Cilliers, J.J., 2017. A flotation control system to optimise performance using peak air recovery. Chemical Engineering Research and Design 117, 57–65. https://doi.org/10.1016/j.cherd.2016.10.021
Shean, B.J., Cilliers, J.J., 2011. A review of froth flotation control. Int J Miner Process 100, 57–71. https://doi.org/10.1016/j.minpro.2011.05.002
Vinnett, L., Marion, C., Grammatikopoulos, T., Waters, K.E., 2020. Analysis of flotation rate distributions to assess erratic performances from size-by-size kinetic tests. Miner Eng 149. https://doi.org/10.1016/j.mineng.2020.106229
Wang, G., Ge, L., Mitra, S., Evans, G.M., Joshi, J.B., Chen, S., 2018. A review of CFD modelling studies on the flotation process. Miner Eng. https://doi.org/10.1016/j.mineng.2018.08.019

Wei, T., Peng, Y., Farrokhpay, S., 2014. Froth stability of coal flotation in saline water. Mineral Processing and Extractive Metallurgy 123, 234–240. https://doi.org/10.1179/1743285514Y.0000000077
Wepener, D.A., le Roux, J.D., Craig, I.K., 2023. Extremum seeking control to optimize mineral recovery of a flotation circuit using peak air recovery. J Process Control 129. https://doi.org/10.1016/j.jprocont.2023.103033
Wepener, D.A., Le Roux, J.D., Craig, I.K., 2022. Extremum Seeking Control of a Flotation Circuit using Peak Air Recovery, in: IFAC-PapersOnLine. Elsevier B.V., pp. 61–66. https://doi.org/10.1016/j.ifacol.2022.09.244
Wills, B.A.., Napier-Munn, Tim., 2006. Wills' Mineral Processing Technology : an Introduction to the Practical Aspects of Ore Treatment and Mineral Recovery. Elsevier Science & Technology Books.
Woodburn, E.T., Austin, L.G., Stockton, J.B., 1994. Froth based flotation kinetic model. Chemical Engineering Research & Design 72, 211–226.
Yianatos, J.B., Bergh, L.G., Díaz, F., 2009. Troubleshooting and Diagnostic of Industrial Flotation Cells.
Ziolkowski, L., le Roux, J.D., Craig, I.K., 2022. Extremum seeking control for optimization of an open-loop grinding mill using grind curves. J Process Control 114, 54–70. https://doi.org/10.1016/j.jprocont.2022.04.005
Ziolkowski, L., le Roux, J.D., Craig, I.K., 2021a. Optimizing grinding mill performance using extremum seeking control, in: IFAC-PapersOnLine. Elsevier B.V., pp. 43–48. https://doi.org/10.1016/j.ifacol.2021.10.048
Ziolkowski, L., le Roux, J.D., Craig, I.K., 2021b. Optimizing grinding mill performance using extremum seeking control. IFAC-PapersOnLine 54, 43–48. https://doi.org/10.1016/J.IFACOL.2021.10.048

APPENDIX A

The appendices presented in this section contain supplementary information that supports the methodology, results, and discussions presented in the main body of this dissertation. This appendix includes simulation results, data tables, additional performance plots, and calculations that were not included in the main chapters.

A1. Model Equations

The equations listed in Table A1 summarise the key equations from literature that form the foundation for this study. These equations describe the relationships between air recovery, froth height, superficial gas velocity and other flotation parameters. The remainder of the section presents a step by step derivation of the air recovery model used in this study.

Table A 1: Summary of key equations from literature.

Equation Number	Equation	Reference
1	$\lambda_{\mathrm{air}} = \frac{hf}{J_g}$	(Oosthuizen et al., 2021)
A1	$\frac{\mathrm{d}}{\mathrm{dt}}\mathrm{D}_{\mathrm{BF_k}} = \frac{(\mathrm{K}_{\mathrm{BS}_{Jg}} J_{g_k} + \mathrm{K}_{\mathrm{BS}_\lambda}\lambda_{\mathrm{air_k}} - \mathrm{D}_{\mathrm{BF_k}})}{\lambda_{\mathrm{air_k}}}$	(Oosthuizen et al., 2021)
A2	$\frac{\mathrm{d}}{\mathrm{dt}}\alpha = \frac{\left(\mathrm{K}_{\alpha_{\mathrm{BF}}}\mathrm{D}_{\mathrm{BF}} + \mathrm{K}_{\alpha_{Jg}} J_g - \mathrm{a}_k\right)}{\lambda_{\mathrm{air}}}$	(Oosthuizen et al., 2021)
A3	$\alpha = \mathrm{A} + \mathrm{B}\lambda_{\mathrm{air}} + \mathrm{C}J_g$	(Brooks et al., 2024)

$$\lambda_{\mathrm{air}} = \frac{hf}{J_g} \tag{1}$$

$$\frac{\mathrm{d}}{\mathrm{dt}}\mathrm{D_B} = \frac{(\mathrm{K}_{\mathrm{BS}_{Jg}} J_g + \mathrm{K}_{\mathrm{BS}_\lambda}\lambda_{\mathrm{air}} - (\mathrm{D}_{\mathrm{BF}} - \mathrm{D}_{\mathrm{BF}_{\mathrm{ref}}})}{\lambda_{\mathrm{air}}} \tag{A4}$$

$$\frac{\mathrm{d}}{\mathrm{dt}}\alpha = \frac{\left(\mathrm{K}_{\alpha_{\mathrm{BF}}}\mathrm{D}_{\mathrm{BF}} + \mathrm{K}_{\alpha_{Jg}}J_g - (\alpha - \alpha_{\mathrm{ref}})\right)}{\lambda_{\mathrm{air}}} \quad (A5)$$

Bubble size model

The method used to develop the air recovery model is shown in the table below. The final developed bubble size model is shown below:

$$\frac{\mathrm{d}}{\mathrm{dt}}\mathrm{D}_{\mathrm{B}} = \frac{(\mathrm{K}_{\mathrm{BS}_{Jg}}J_g + \mathrm{K}_{\mathrm{BS}_{\lambda}}\lambda_{\mathrm{air}} - (\mathrm{D}_{\mathrm{BF}} - \mathrm{D}_{\mathrm{BF}_{\mathrm{ref}}})}{\lambda_{\mathrm{air}}} \quad (A4)$$

At steady state:

$$0 = \frac{(\mathrm{K}_{\mathrm{BS}_{Jg}}J_g + \mathrm{K}_{\mathrm{BS}_{\lambda}}\lambda_{\mathrm{air}} - (\mathrm{D}_{\mathrm{BF}} - \mathrm{D}_{\mathrm{BF}_{\mathrm{ref}}})}{\lambda_{\mathrm{air}}} \quad (A6)$$

$$0 = (\mathrm{K}_{\mathrm{BS}_{Jg}}J_g + \mathrm{K}_{\mathrm{BS}_{\lambda}}\lambda_{\mathrm{air}} - (\mathrm{D}_{\mathrm{BF}} - \mathrm{D}_{\mathrm{BF}_{\mathrm{ref}}}) \quad (A7)$$

$$\mathrm{D}_{\mathrm{BF}} = \mathrm{D}_{\mathrm{BF}_{\mathrm{ref}}} + \mathrm{K}_{\mathrm{BS}_{Jg}}J_g + \mathrm{K}_{\mathrm{BS}_{\lambda}}\lambda_{\mathrm{air}} \quad (A8)$$

Substituting Equation A8 into Equation A5

From Equation A5:

$$\frac{\mathrm{d}}{\mathrm{dt}}\alpha = \frac{\left(\mathrm{K}_{\alpha_{\mathrm{BF}}}\mathrm{D}_{\mathrm{BF}} + \mathrm{K}_{\alpha_{Jg}}J_g - (\alpha - \alpha_{\mathrm{ref}})\right)}{\lambda_{\mathrm{air}}} \quad (A5)$$

$$\frac{\mathrm{d}}{\mathrm{dt}}\alpha = \frac{\left(\mathrm{K}_{\alpha_{\mathrm{BF}}}\left(\mathrm{D}_{\mathrm{BF}_{\mathrm{ref}}} + \mathrm{K}_{\mathrm{BS}_{Jg}}J_g + \mathrm{K}_{\mathrm{BS}_{\lambda}}\lambda_{\mathrm{air}}\right) + \mathrm{K}_{\alpha_{Jg}}J_g - (\alpha - \alpha_{\mathrm{ref}})\right)}{\lambda_{\mathrm{air}}} \quad (A9)$$

$$\frac{\mathrm{d}}{\mathrm{dt}}\alpha = \frac{\left(\mathrm{K}_{\alpha_{\mathrm{BF}}}\mathrm{D}_{\mathrm{BF}_{\mathrm{ref}}} + \mathrm{K}_{\alpha_{\mathrm{BF}}}\mathrm{K}_{\mathrm{BS}_{Jg}}J_g + \mathrm{K}_{\alpha_{\mathrm{BF}}}\mathrm{K}_{\mathrm{BS}_{\lambda}}\lambda_{\mathrm{air}} + \mathrm{K}_{\alpha_{Jg}}J_g - (\alpha - \alpha_{\mathrm{ref}})\right)}{\lambda_{\mathrm{air}}} \quad (A10)$$

$$\frac{\mathrm{d}}{\mathrm{dt}}\alpha = \frac{\mathrm{K}_{\alpha_{\mathrm{BF}}}\mathrm{D}_{\mathrm{BF}_{\mathrm{ref}}} + J_g(\mathrm{K}_{\alpha_{\mathrm{BF}}}\mathrm{K}_{\mathrm{BS}_{Jg}} + \mathrm{K}_{\alpha_{Jg}}) + \mathrm{K}_{\alpha_{\mathrm{BF}}}\mathrm{K}_{\mathrm{BS}_\lambda}\lambda_{\mathrm{air}} - (\alpha - \alpha_{\mathrm{ref}})}{\lambda_{\mathrm{air}}} \quad (A11)$$

At steady state:

$$0 = \frac{\mathrm{K}_{\alpha_{\mathrm{BF}}}\mathrm{D}_{\mathrm{BF}_{\mathrm{ref}}} + J_g(\mathrm{K}_{\alpha_{\mathrm{BF}}}\mathrm{K}_{\mathrm{BS}_{Jg}} + \mathrm{K}_{\alpha_{Jg}}) + \mathrm{K}_{\alpha_{\mathrm{BF}}}\mathrm{K}_{\mathrm{BS}_\lambda}\lambda_{\mathrm{air}} - (\alpha - \alpha_{\mathrm{ref}})}{\lambda_{\mathrm{air}}} \quad (A12)$$

$$\alpha - \alpha_{\mathrm{ref}} = \mathrm{K}_{\alpha_{\mathrm{BF}}}\mathrm{D}_{\mathrm{BF}_{\mathrm{ref}}} + J_g(\mathrm{K}_{\alpha_{\mathrm{BF}}}\mathrm{K}_{\mathrm{BS}_{Jg}} + \mathrm{K}_{\alpha_{Jg}}) + \mathrm{K}_{\alpha_{\mathrm{BF}}}\mathrm{K}_{\mathrm{BS}_\lambda}\lambda_{\mathrm{air}} \quad (A13)$$

$$\alpha = \alpha_{\mathrm{ref}} + \mathrm{K}_{\alpha_{\mathrm{BF}}}\mathrm{D}_{\mathrm{BF}_{\mathrm{ref}}} + J_g(\mathrm{K}_{\alpha_{\mathrm{BF}}}\mathrm{K}_{\mathrm{BS}_{Jg}} + \mathrm{K}_{\alpha_{Jg}}) + \mathrm{K}_{\alpha_{\mathrm{BF}}}\mathrm{K}_{\mathrm{BS}_\lambda}\lambda_{\mathrm{air}} \quad (A14)$$

$$\alpha = (\alpha_{\mathrm{ref}} + \mathrm{K}_{\alpha_{\mathrm{BF}}}\mathrm{D}_{\mathrm{BF}_{\mathrm{ref}}}) + \mathrm{K}_{\alpha_{\mathrm{BF}}}\mathrm{K}_{\mathrm{BS}_\lambda}\lambda_{\mathrm{air}} + J_g\left(\mathrm{K}_{\alpha_{\mathrm{BF}}}\mathrm{K}_{BS_{Jg}} + \mathrm{K}_{\alpha_{Jg}}\right) \quad (A15)$$

or

$$\alpha = \mathrm{A} + \mathrm{B}\lambda_{\mathrm{air}} + \mathrm{C}J_g \quad (A3)$$

$$\alpha = \mathrm{A} + \mathrm{B}\frac{hf}{J_g} + \mathrm{C}J_g \quad (9)$$

A2. Filter configurations

This section presents the detailed performance metrics for the various proportional gain values tested including the exponential smoothing filter configurations. The tables report the performance metrics for air recovery and superficial gas velocity. The results are discussed in the Chapter 4 and are used to evaluate the trade offs between noise suppression, responsiveness, and overall controller stability for each configuration.

Table A 2: Performance metrics for air recovery at $K_p = 0.5$

$K_p = 0.5$	Filter configuration				
Metric	$G_1 = 02$, $G_2 = 0.8$	$G_1 = 0.4$, $G_2 = 0.6$	$G_1, G_2 = 0.5$	$G_1 = 0.6$, $G_2 = 0.4$	$G_1 = 0.8$, $G_2 = 0.2$
MAE	0.01	0.01	0.03	0.01	0.01
SD	0.01	0.01	0.02	0.02	0.01
RMSE	0.00	0.01	0.02	0.01	0.01
IAE	0.72	0.76	2.56	1.18	1.50
CPI	97.63	93.29	27.76	59.94	47.26

Table A 3: Performance metrics for superficial gas velocity at $K_p = 0.5$

$K_p = 0.5$	Filter configuration				
Metric	$G_1 = 02$, $G_2 = 0.8$	$G_1 = 0.4$, $G_2 = 0.6$	$G_1, G_2 = 0.5$	$G_1 = 0.6$, $G_2 = 0.4$	$G_1 = 0.8$, $G_2 = 0.2$
MAE	0.36	0.29	0.55	0.29	0.39
SD	0.54	0.30	0.40	0.32	0.26
RMSE	0.26	0.21	0.52	0.16	0.38
IAE	36.43	28.94	55.88	29.13	39.38
CPI	20.80	26.18	13.56	26.00	19.23

Table A 4: Performance metrics for air recovery at $K_p = 1.5$

$K_p = 1.5$	Filter configuration				
Metric	$G_1 = 02$, $G_2 = 0.8$	$G_1 = 0.4$, $G_2 = 0.6$	$G_1, G_2 = 0.5$	$G_1 = 0.6$, $G_2 = 0.4$	$G_1 = 0.8$, $G_2 = 0.2$
MAE	0.02	0.01	0.01	0.02	0.01
SD	0.02	0.01	0.01	0.02	0.02
RMSE	0.02	0.01	0.01	0.02	0.01
IAE	2.31	0.91	1.10	1.70	1.44
CPI	30.61	77.64	64.53	41.67	49.00

Table A 5: Performance metrics for superficial gas velocity at $K_p = 1.5$

$K_p = 1.5$	Filter configuration				
Metric	$G_1 = 02$, $G_2 = 0.8$	$G_1 = 0.4$, $G_2 = 0.6$	$G_1, G_2 = 0.5$	$G_1 = 0.6$, $G_2 = 0.4$	$G_1 = 0.8$, $G_2 = 0.2$
MAE	0.36	0.36	0.39	0.48	0.43
SD	0.53	0.40	0.46	0.52	0.46
RMSE	0.07	0.21	0.14	0.44	0.29
IAE	36.37	36.45	39.52	48.44	43.17
CPI	20.83	20.78	19.17	15.64	17.55

A3. Cost Function Analysis Results

This section provides the grid search sweep cost function analysis results discussed in Chapter 4.

Table A 6: Sweep cost function analysis results.

Proportional gain (K_p)	SSE	SD
0.20	2.92E-04	1.34E-02
0.25	3.28E-04	9.66E-03
0.30	3.23E-04	8.31E-03
0.35	2.96E-04	9.75E-03
0.40	2.78E-04	1.48E-02
0.45	3.05E-04	1.14E-02
0.50	2.89E-04	4.47E-03
0.55	3.45E-04	6.85E-03
0.60	2.89E-04	1.86E-03
0.65	2.86E-04	1.98E-03
0.70	2.80E-04	9.90E-04
0.75	2.78E-04	4.33E-03
0.80	2.95E-04	2.08E-03
0.85	2.73E-04	1.12E-03
0.90	3.12E-04	1.06E-03
0.95	2.89E-04	1.68E-03
1.00	3.08E-04	1.18E-03
1.05	3.07E-04	3.35E-03
1.10	2.99E-04	1.25E-03
1.15	3.04E-04	2.06E-03
1.20	3.56E-04	2.31E-03
1.25	3.65E-04	9.52E-04
1.30	3.01E-04	1.40E-03
1.35	2.85E-04	1.07E-03
1.40	3.07E-04	1.42E-03
1.45	3.43E-04	2.13E-03
1.50	3.07E-04	1.19E-03

A4. Analytical Method

This section expands on the analytical approach outlined in Chapter 3, providing full derivations, methodology, results and discussions. In this model, K is a mathematical proportionality constant that links the gradient $\frac{\mathrm{d}\alpha}{\mathrm{dJ_g}}$ to changes in superficial gas velocity. It differs from the proportional gain (Kp) used in the main simulation-based controller tuning. Positive K values represent corrective feedback driving the system toward PAR, while negative values illustrate the instability that results from applying feedback in the wrong direction.

A4.1 Analytical Approach Methodology

The fundamental equation describing the system dynamics is expressed as:

$$\frac{\mathrm{d}J_g}{\mathrm{dt}} = \mathrm{K}\frac{\mathrm{d}\alpha}{\mathrm{d}J_g} \qquad (\mathrm{A}\ 16)$$

where K is a proportionality constant governing the rate at which the system converges to PAR.

The air recovery function is given by:

$$\alpha = \mathrm{A} + \mathrm{B}\frac{h_f}{J_g} + \mathrm{C}J_g \qquad (8)$$

where A, B, and C are empirical constants characterising the system.

Differentiating α with respect to J_g gives:

$$\frac{\mathrm{d}\alpha}{\mathrm{dJ_g}} = -\frac{\mathrm{B}_{h_f}}{{J_g}^2} + \mathrm{C} \qquad (16)$$

Substituting this expression into the governing equation for J_g (Equation A 16) :

$$\frac{\mathrm{dJ_g}}{\mathrm{dt}} = \mathrm{K}\left(\mathrm{C} - \frac{\mathrm{B}_{h_f}}{{J_g}^2}\right) \qquad (\mathrm{A}\ 17)$$

Defining the optimal J_g as:

$$J_{g_{opt}} = \sqrt{\frac{Bh_f}{C}} \qquad (A\ 18)$$

the governing equation became:

$$\frac{dJ_g}{dt} = KC\left(1\ \frac{J_{g^2{}_{opt}}}{J_{g^2}}\right) = KC\ \frac{J_{g^2} - J_{g^2{}_{opt}}}{J_{g^2}} \qquad (A\ 19)$$

separating the variables:

$$\left(\frac{J_{g^2}}{J_{g^2} - J_{g^2{}_{opt}}}\right) dJ_g = \mathrm{CK}\, dt \qquad (\mathrm{A}\ 20)$$

integrating both sides gives the implicit solution:

$$\mathrm{t} = -\frac{1}{\mathrm{CK}}\left[J_g - J_{g_0} + \frac{J_{g_{opt}}}{2} ln\left(\frac{(J_g - J_{g_{opt}})(J_{g_0} + J_{g_{opt}})}{(J_g + J_{g_{opt}})(J_{g_0} - J_{g_{opt}})}\right)\right] \qquad (24)$$

Where J_{g_0} is the initial superficial gas velocity at t = 0. This equation describes the system's time response as it approaches $J_{g_{opt}}$. To facilitate this integration, it was assumed that the differential equation could be rearranged in the form $\left(J_g - J_{g_{opt}}\right)\frac{dJ_g}{dt} = CK$, which simplifies the nonlinear behaviour near PAR and enables an implicit solution. This solution was computed iteratively by incrementing J_g in small steps and evaluating t, α, and $\frac{\mathrm{d}\alpha}{\mathrm{d}J_g}$ at each step. The values were then plotted to visualize the trend of α as a function of J_g, and the results are presented in the appendix A.

A4.2 Modelling Air Recovery using a Quadratic model

An alternative analytical was developed by approximating the air recovery function with a quadratic expression:

$$\alpha = \mathrm{D} + \mathrm{E}J_g + \mathrm{F}{J_g}^2 \quad (25)$$

where D, E, and F are coefficients determined through empirical fitting.

differentiating α with respect to J_g gave:

$$\frac{\mathrm{d}\alpha}{\mathrm{d}J_g} = \mathrm{E} + 2\mathrm{F}J_g \quad (26)$$

setting this to zero at PAR yielded:

$$J_{g_{opt}} = -\frac{\mathrm{E}}{2\mathrm{F}} \quad (27)$$

substituting into the main ODE:

$$\frac{\mathrm{d}J_g}{\mathrm{dt}} = \mathrm{K}\left(\mathrm{E} + 2\mathrm{F}J_g\right) = 2\mathrm{KF}\ (J_g - J_{g_{opt}}) \quad (28)$$

which has the solution:

$$J_g(t) = J_{g_{opt}} + \left[J_{g_0} - J_{g_{opt}}\right] \mathrm{e}^{2\mathrm{FKt}} \quad (29)$$

where J_{g_0} is the initial condition at t = 0.

The time constant was given as:

$$\tau = \frac{1}{2\mathrm{FK}} \quad (30)$$

This time constant defines how quickly the system responds to changes in superficial gas velocity. The coefficients D, E, and F were determined in two ways: first, through curve fitting to match the previously computed air recovery function, and second, by performing a second-order Taylor expansion around $J_{g_{opt}}$, yielding:

For the Taylor expansion of the original α about $J_{g_{opt}}$:

$$D = A + 3CJ_{g_{opt}} \quad (31)$$

$$E = -2C \quad (32)$$

$$F = \frac{C}{J_{g_{opt}}} \quad (33)$$

This formulation ensures that the quadratic approximation aligns with the maximum air recovery at $J_{g_{opt}}$. The quadratic model was plotted alongside the original function to compare their accuracy in representing the system behaviour near PAR.

A5. Analytical Method Results and Discussion

The analytical solution provided a mathematical basis for understanding the relationship between air recovery and superficial gas velocity. The results presented in Figures A1 to A11 illustrate the system's behaviour under both ideal and disturbed conditions. In each case, the analytical results were compared against the simulation model used in the main study to assess accuracy, the influence of the proportional gain on performance, and the implications for the control strategy developed in this study. The parameters used for the analysis are summarised in Table A 7.

Table A 7: Parameters used in the analytical solution for various K values.

A	4.92
B	-15.00
C	-0.28
J_{g_0}	5.00
K_p	0.50; 1.00; 1.50
$J_{g_{opt}}$	7.54
h_f	106 mm
Noise level (σ)	0.10; 0.20

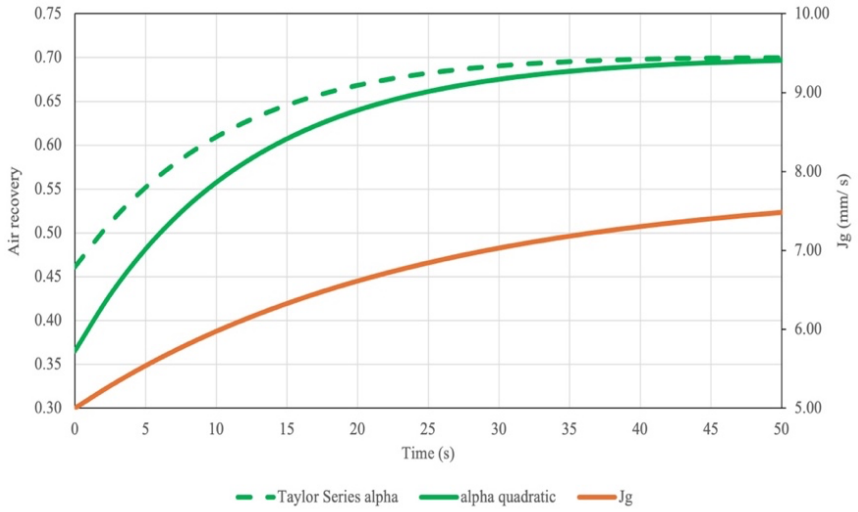


Figure A 1: Comparison of the analytical model results for α and J_g at K = 0.5.

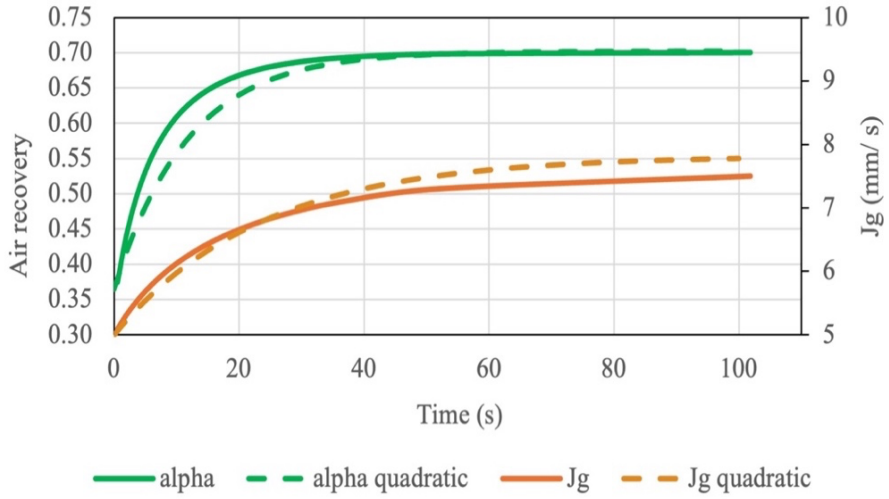


Figure A 2:Graphical representation of α and J_g over time for K = 0.5.

In the absence of external disturbance, both the quadratic and Taylor expansion series analytical models closely followed the baseline air recovery from the simulation. The quadratic approximation successfully captured the non-linear relationship between the superficial gas velocity and air recovery, accurately reproducing the expected peak and stabilisation at the point of PAR. The Taylor series analytical model also produced results that aligned with the simulation in the area of PAR as shown in Figures A1 and A2.

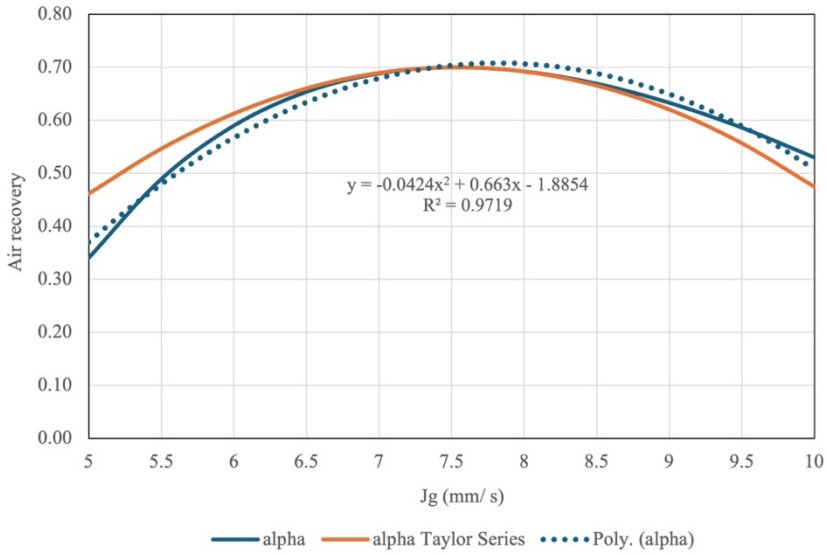


Figure A 3: α calculated using the Taylor Series and the polynomial approximations against J_g at K = 1

Figure A3 further highlights the accuracy of the analytical approach, showing the fitted polynomial trend line against the simulation derived air recovery values. The high coefficient of determination ($R^2 = 0.972$) confirms that the model explains most of the variance in the data, accurately showing the location of PAR. This alignment indicates that in the absence of disturbances, the analytical methods provide a reliable estimate of both the optimal superficial gas velocity and the curve towards PAR.

Increasing K in these ideal conditions from 0.5 to 1.0 improved the convergence rate to PAR without introducing instability under these conditions. Notably, at K = 1 the system converged on PAR significantly faster, while lower gains such as K = 0.5 approached PAR more gradually over an extended period.

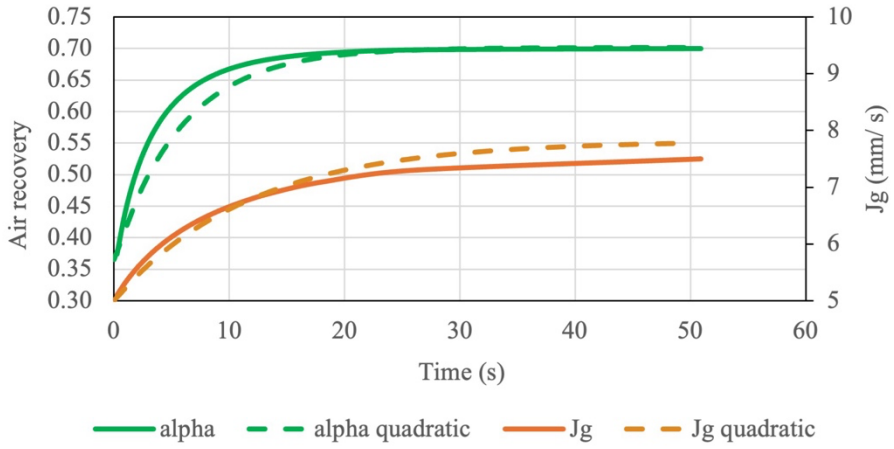


Figure A 4: Comparison of α and J_g over time, including both the calculated values and their quadratic approximations at K = 1.

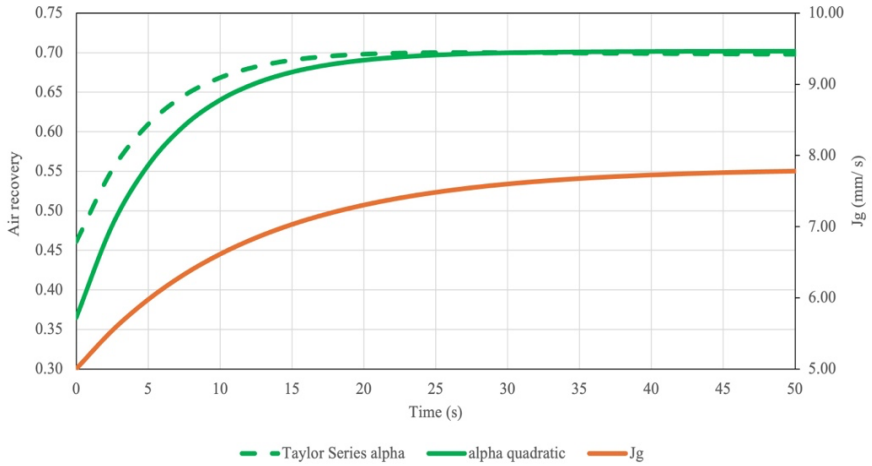


Figure A 5: Evolution of air recovery and J_g over time at K = 1.

When disturbances were introduced as shown in Figures A6 to A, the differences between the analytical and simulation models became more noticeable. At the lowest tested noise level, the quadratic approach tracked the simulation close to PAR, with K = 0.5 producing smaller fluctuations compared to K = 1.0, which responded more aggressively to noise. As noise increased, fluctuations in both air recovery and superficial gas velocity lasted for longer periods at higher gains, delaying stable convergence. At the highest noise level ($\sigma = 0.2$), both methods showed significant variability, and the analytical models could only reproduce the average trend rather than the rapid fluctuations observed in the simulation.

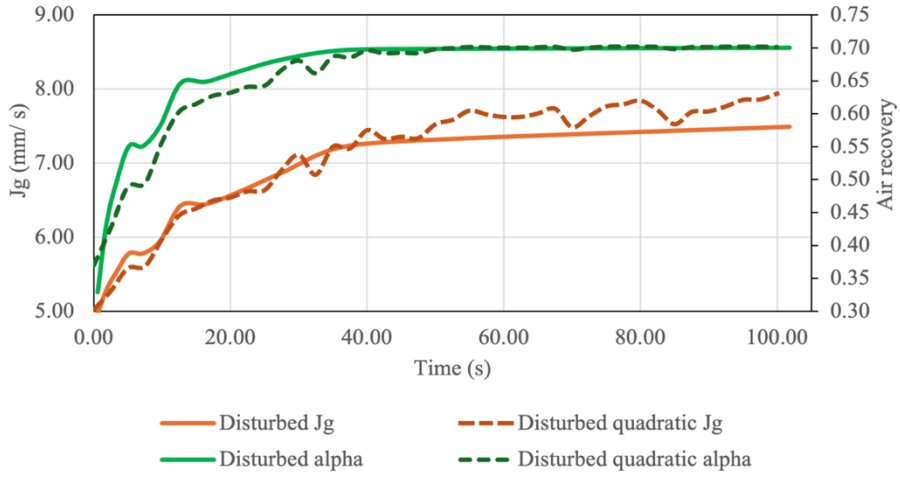


Figure A 6: Disturbed J_g, α and quadratic approximations with a σ of 0.1 with K = 0.5.

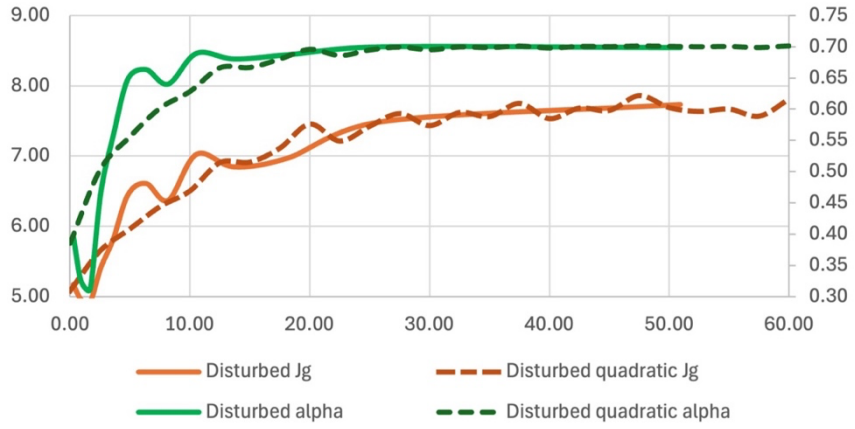


Figure A 7: Disturbed J_g and α over time with σ = 0.1 and K= 1.

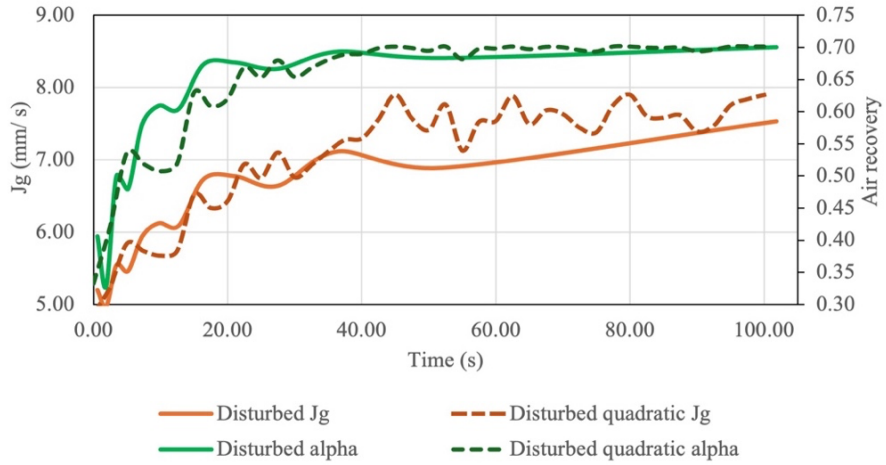


Figure A 8: J_g and α with a noise level of; σ = 0.2 and K =0.5.

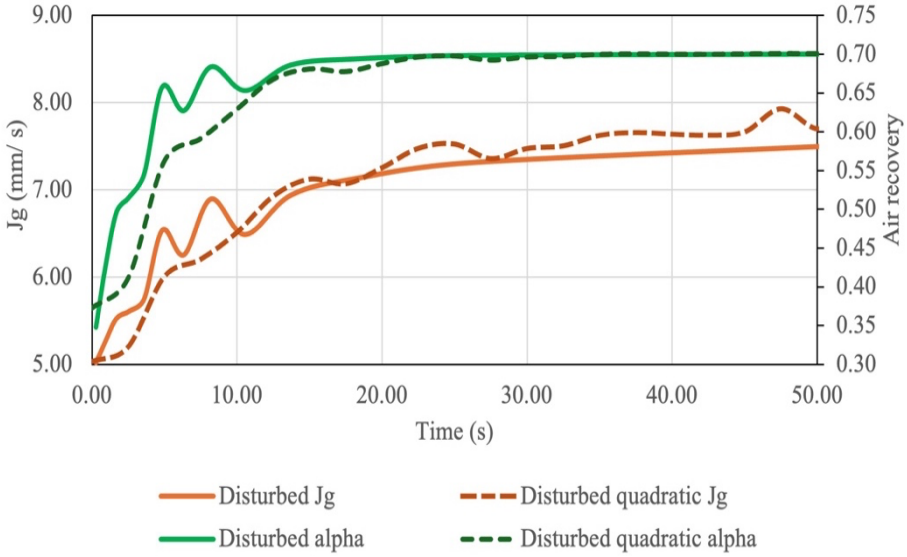


Figure A 9: Disturbed α, J_g and disturbed quadratic approximations with σ of 0.2 with K 1.

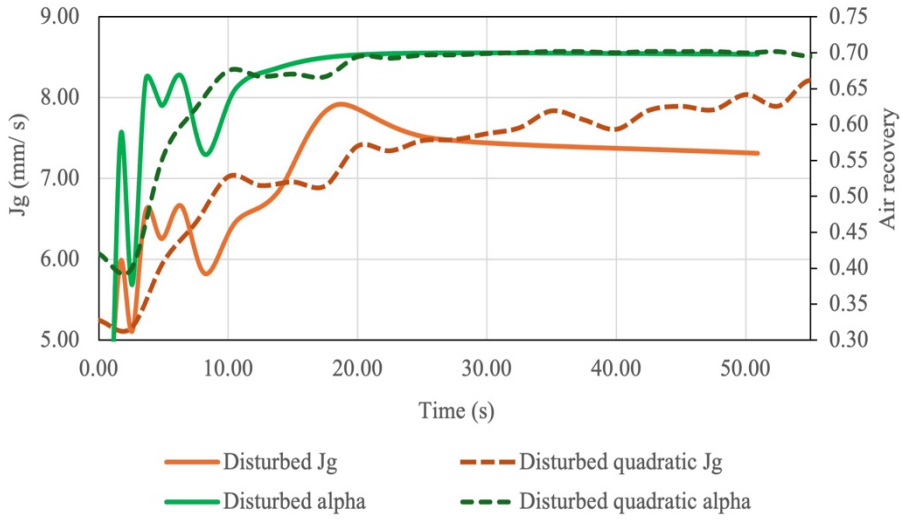


Figure A 10: Disturbed J_g , α and quadratic approximations with a σ of 0.2 and K= 1.

Across all scenarios, the results confirmed that the analytical methods were effective in predicting the position of PAR and the general shape of the approach trajectory in low disturbance conditions. However, their inability to fully account for short term process variability under higher noise levels and disturbances highlighted the importance of real time adaptability in practical applications. These findings support the simulation results in the main body of the dissertation, demonstrating that while higher gains can reduce convergence time under ideal conditions, in noisy environments a balance between responsiveness and stability must be maintained.