Accepted Manuscript

Title: A FLOTATION CONTROL SYSTEM TO OPTIMISE PERFORMANCE USING PEAK AIR RECOVERY

Author: B. Shean K. Hadler J.J. Cilliers

Please cite this article as: Shean, B., Hadler, K., Cilliers, J.J., A FLOTATION CONTROL SYSTEM TO OPTIMISE PERFORMANCE USING PEAK AIR RECOVERY.Chemical Engineering Research and Design http://dx.doi.org/10.1016/j.cherd.2016.10.021
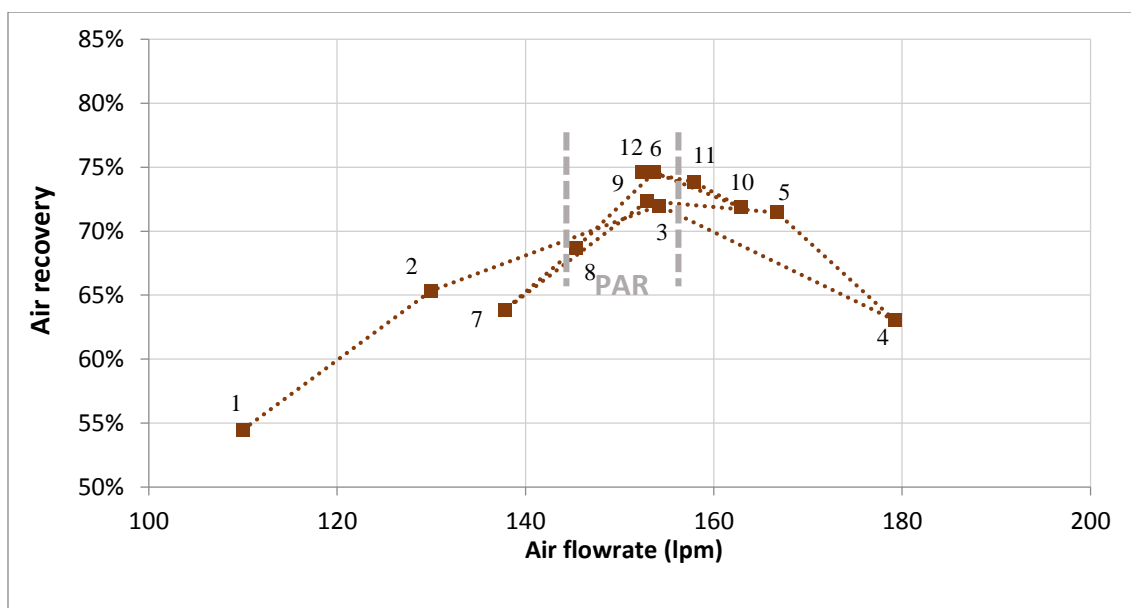
A FLOTATION CONTROL SYSTEM TO OPTIMISE PERFORMANCE USING PEAK AIR RECOVERY

B. Shean, K. Hadler*, J.J. Cilliers

Department of Earth Science and Engineering, Imperial College London, SW7 2AZ ,UK

*Corresponding Author

Graphical Highlight

Expected air recovery as a function of air rate (heavy dashed line) and air recoveries obtained using the GSS-based PAR seeking control system starting from an air rate lower than the PAR air rate. The order in which the air rates were selected by the control system is given by the numbered data points

Research Highlights

- Novel automated air rate control for froth flotation based on Peak Air Recovery
- 70 l continuous laboratory flotation system used to implement the control system
- The proposed control system was based on the direct search algorithm
- Flotation cell controlled to the PAR air rate both at steady and unsteady state

ABSTRACT

Automatic control of industrial flotation cells and circuits presents a set of significant challenges due to the number of variables, the sensitivity of flotation cells to variation in these variables and the complexity of predicting flotation performance and/or developing a strategy for optimisation. Air recovery, a measure of froth stability, has been shown to pass through a peak as flotation cell aeration increases. Furthermore, the air rate at which the Peak Air Recovery (PAR) is obtained results in optimal flotation performance, whether improved concentrate grade, recovery or both grade and recovery. Peak air recovery, therefore, presents a clear optimising control strategy for the operation of flotation cells which is generic to all flotation cells regardless of position in the flotation circuit. In this study, a novel control system based on PAR is developed and demonstrated using a large continuous laboratory flotation cell.

In this study, a direct search optimisation algorithm based on the GSS (generating set search) methodology was developed using a 70 l continuous flotation cell operating with a two-phase system (surfactant solution and air only). Characterisation of the laboratory system showed that it was stable for up to 6 hours and exhibited a

2

reproducible peak in air recovery. A dynamic model of the response of the system with regards to changes in air recovery was developed that allowed simulations of the proposed optimising control system to be carried out. The optimisation algorithm was then applied to the experimental system. The trialled GSS algorithm was shown to find the PAR air rate when starting above, below and at the PAR air rate, and additionally with a disturbance introduced into the system. While the direct search approach can be slow, it is simple and robust. This demonstrates an innovative approach to optimising control for froth flotation and is the first application of froth stability maximisation for flotation control.

1. INTRODUCTION

Froth flotation is a well-established and widely applied separation process, however the separation performance is dependent on numerous variables, including feed grade, particle size distribution, mineralogy (of both valuable minerals and gangue), slurry solids content, reagent type and addition and operating conditions such as air rate and froth depth. Variation in any one of these aspects results in changes in the chemical conditions and physical sub-processes in the pulp phase, which consequently affects the froth structure and stability and therefore the concentrate grade and the recovery of valuable minerals. The complex nature of flotation operation means that control of industrial flotation cells continues to present a significant set of challenges (Bergh and Yianatos, 2011), but the potential rewards in terms of improved metallurgical performance are substantial (McKee, 1991).

A comprehensive review of flotation control can be found in Shean and Cilliers (2011), however in general, flotation control occurs at several levels, from regulatory control to maintain primary variables at setpoints (e.g. pulp level) to optimising control to maximise profitability, as described by Hodouin et al. (2001) and Laurila et al. (2002). Advanced or optimising control has been demonstrated through the application of model predictive control (Maldonado et al., 2009; Bergh and Yianatos, 2011) and expert systems (Aldrich et al., 1997).

Model predictive control (MPC) offers the potential for high quality advanced control of flotation cells, however it relies on dynamic models accounting for the relationship between operating parameters and measured responses (Bergh and Yianatos, 2011). The modelling of flotation systems, however, remains a much-studied field of research, not least due to the complex interactions of all variables (Mathe et al., 1998; Wang et al., 2015). To this end, much research has been focused on expert and fuzzy control systems for flotation cells, at both lab (Jahedsaravani et al., 2016) and plant scale (Bergh et al., 1999; Osorio et al., 1999; Nunez et al., 2010), although hybrid MPC systems have also been suggested (Karelovic et al., 2015; Putz and Cipriano, 2015). The development and advancement of machine vision technology, in particular, has enabled control systems based on fuzzy control to become more widely adopted (Aldrich et al., 2010).

Machine vision quantifies the visual attributes of the froth, such as bubble size, froth velocity, colour and bubble stability which are often used for manual flotation control by experienced operators. All of these features have been used as the basis for flotation control systems in recent years, as reviewed by Aldrich et al. (2010). For example Zhu et al. (2014) developed a predictive controller for a copper flotation plant based on manipulating reagent dosage to obtain a specified bubble size distribution, with feed grade as an additional input. Kaartinen et al. (2006) used an image analysis system to develop an expert controller for a zinc flotation plant, using image variables in feedback control to tailor reagent dosage. This resulted in an estimated 1.3% increase in recovery. Supomo et al. (2008), on the other hand, used a commercial image analysis system to control solids flowrate to concentrate down a bank of flotation cells using froth depth, demonstrating improved plant stability and recovery. In all of these systems, however, some 'tuning' of the controller is required in order to establish an acceptable range of the target. A generic optimising control strategy applicable to all flotation cells with a single objective function has, however, not yet been demonstrated.

Air recovery is a measure of froth stability and is defined as the fraction of air entering a flotation cell that overflows the cell lip as unburst bubbles. It has been shown that air recovery passes through a peak as the cell air rate increases and that this corresponds to the air rate at which the highest mineral recovery is obtained (Hadler and Cilliers, 2009). This is shown schematically in Figure 1.

Improved performance through air recovery optimisation, while conceptually simple, can be explained by considering the mechanisms by which particles (both valuable and gangue) are recovered to the concentrate and by the flow behaviour of the overflowing froth (Hadler and Cilliers, 2009). It is measured using the overflowing froth velocity, the height above the cell lip of the overflowing froth and the inlet air flowrate and offers a practical control strategy based on fundamental froth behaviour. Moreover, it allows a single objective function to be used in the optimisation of both concentrate grade and recovery, since altering cell aeration to operate at the peak air recovery air rate can lead to both higher grade and recovery in the case where cells are over-aerated (Hadler et al., 2010).

While the existence of a peak in air recovery has been demonstrated for several industrial and laboratory systems (Hadler and Cilliers, 2009; Hadler et al., 2010; Smith et al., 2010; Qu et al., 2013), the PAR air rate varies with operating variables such as froth depth (Hadler et al., 2012), particle size (Norori-McCormac et al., 2014) and reagent addition (Qu et al., 2013). Neethling and Cilliers (2008) demonstrated the existence of the peak in air recovery using fundamental physics-based modelling of the froth, however there are currently no predictive models available to determine the PAR air rate as a function of the numerous variables in industrial flotation systems. In order, therefore, to develop an optimising control system based on PAR, algorithms are required that do not require models of the system.

One such group of optimisation methods are the direct search methods, which are used when reliable models for a given system are unavailable. Direct search methods focus on finding the 'direction' in which a system must move, and the step sizes that must be taken, such that some 'profit function' is maximised as rapidly as possible (Marlin, 2000). While they can be slow to converge (Kolda et al., 2003), their ease of use make them applicable for process control where models are overly complex or unavailable and where there is significant noise (Xiong and Jutan, 2003). This makes them particularly suited to a flotation control system based on the PAR concept. Direct search methods have been used in flotation circuit optimisation (Dey et al., 1989), however there are no reported studies of direct search optimisation for flotation cell control.

The term 'direct search optimisation' encompasses a group of possible methods, with four of the most commonly mentioned methods being the Generating Set Search (GSS) method, the Golden Section Search method (GLDSS), the Inverse Parabolic Interpolation (IPI) search method, and the Function Fitting Search (FFS) method (Kolda et al., 2003). GSS is slow but robust and is capable of multidimensional, constrained optimisation, widely regarded as the most robust method (Kolda et al., 2003). A simplified example of the GSS method in 2 dimensions is illustrated in Figure 2.

The GSS algorithm starts by computing the unknown objective function $f(x)$ at an initial point, as in Figure 2a. The algorithm then takes steps of equal size of the independent variable x until $f(x)$ no longer increases (Figure 2a – c). This is an indication that the maximum has been bracketed (i.e. the maximum lies between the initial value of x and the value of x in Figure 2c). At this stage the step size is decreased to allow for convergence and the search direction is reversed. This is repeated until the step size falls below a minimum tolerance, at which stage the system is considered to have converged.

This study demonstrates the first application of an optimising control system based on air rate control to find and maintain peak air recovery using a large laboratory scale flotation cell. The flotation cell operates with a two-phase surfactant-air system (i.e. there is no separation of solids) exhibiting froth behaviour analogous to industrial flotation operations. Following characterisation of the air recovery response of the laboratory system, a dynamic

simulator was developed in order to test the efficacy of the control algorithm, before experimental trials were undertaken.

2. EXPERIMENTAL SYSTEM

Flotation System

All tests were carried out using a custom built 70 l flotation cell (height 50 cm) with a sloped launder attached to the outside of the tank to allow removal of the overflowing froth. The tank contained four baffles in the pulp phase to promote mixing and agitation was achieved using an overhead mixer (rated 0.25 kW) with six blade Rushton impeller operated at 1100 rpm. The air was introduced into the cell directly below the impeller via tubing anchored to the base of the tank, with an elbow joint to direct the air flow upwards into the shear zone created by the impeller. A diagram of the cell is given in Figure 3. This system generated sufficient shear to create bubbles of approximately 1 mm (estimated visually) in the pulp zone, with a stable pulp-froth interface over a range of air rates. The flotation cell operated continuously in a closed loop with a reservoir (capacity 50 l) collecting the overflowing froth from the launder and providing feed to the cell.

The flotation system tested in this study was based on a two-phase system (i.e. surfactant solution and air), with no separation of solid particles carried out. This allowed for continuous operation of the cell over long time periods in order to test the control system more rigorously. To generate a froth phase for this two-phase system which exhibited behaviour analogous to that of typical industrial flotation separations, the following criteria were targeted:

- The froth should be mobile and flow towards the cell lip
- The froth should be sufficiently stable to overflow the launder and break down once in the launder
- Coalescence of bubbles in the froth phase should occur, with a distinct range of bubble size present at the froth surface
- The froth surface should be opaque with contrast between individual bubbles in order to enable adequate image analysis for the velocity and overflowing height measurements

To achieve the required froth properties, a polyglycol frother (DowFroth 250) was used at a concentration of 7.5 ppm and xanthan gum was added as a viscosity modifier at 70 ppm (Brito-Parada and Cilliers, 2012). The addition of a viscosity modifier hindered drainage in the froth phase, where unattached particles would increase the viscosity of the slurry in the Plateau borders in industrial froths (Pugh, 1996; Neethling and Cilliers, 2003). While this produced a coalescing overflowing froth, there was insufficient contrast between bubbles at the froth surface, resulting in erroneous image analysis data. To address this, ultramarine blue U8 dye pigment was added at 1000 ppm (d_{50} approximately 8 μm) with the cationic surfactant dodecyl trimethylammonium bromide, DTAB, added at 0.7 ppm to impart further hydrophobicity to the particles, in addition to stabilising the froth. Despite the addition of ultramarine pigment, the system was still considered to be two phase, as the proportion of solid particles was low and intended to provide colour only.

Further addition of frother was required throughout the tests due to a reduction in froth overflowing the lip over the test duration following an initial 60 minute period of stability. The additional frother was added to the reservoir tank at low flowrates for all tests and an initial period of 60 minutes 'stabilisation time' was allowed. This system was demonstrated to be stable for up to 6 hours.

Finally, as the system was two phase, it was found to be sensitive to environmental humidity, as demonstrated previously by Li et al. (2010). The laboratory relative humidity was observed to vary between 20 and 80% depending on weather and season, affecting significantly the behaviour of the froth and producing poor experimental repeatability. To control the effects of humidity, an enclosure was constructed around the entire system, allowing the relative humidity to be maintained at between 70 and 75% and the temperature to be regulated between 17 and 20 ^{0}C.

Air Recovery Measurement System

Air recovery is determined using the following relationship (Woodburn et al., 1994):

$$\alpha = \frac{v_f h_f L}{Q_a} \qquad\qquad \text{Equation 1}$$

Where α is the air recovery, v_f is the overflowing froth velocity (in cm/s), h_f is the overflowing froth height (in cm), L is the overflowing lip length (cm) and Q_a is the volumetric air flowrate entering the cell (cm³/s). This relationship assumes that the liquid content of the overflowing froth is negligible, which is acceptable for operation of a flotation cell where the tank is not 'sliming' (i.e. under conditions where there is loss of the interface and the pulp overflows the lip).

Air rate to the flotation cell was controlled using a mass flow controller with an operating range of 0 – 200 lpm for a supply pressure of 2 – 2.5 bar (gauge). Experiments were carried out in the range of 130 – 190 lpm, corresponding to a superficial gas velocity (J_g) range of 1.10 – 1.61 cm/s. A maximum change of 80 lpm was achieved in a time of 13.6 s, however in the development of the control system, the maximum allowable air rate change was 25 lpm, which could be achieved in 5 s.

The overflowing froth velocity was measured using an in-house developed block matching image analysis algorithm that was interfaced with LabVIEW. Development of image analysis techniques for the measurement of froth surface properties has been the focus of much research in the past 20 years; a comprehensive review is given in Aldrich et al. (2010). The block matching algorithm is one of several established techniques for tracking bubbles at the froth surface. The conversion of the froth velocity from pix/fr to cm/s accounted for the variation in froth surface height from the camera lens.

The overflowing froth height was measured by optical level sensor, based on a visible laser pulse. The laser used for this system (IFM electronic) was designed for the measurement of tank/silo levels, with an operating range of 0.2 – 10 m. The laser measured the height to the froth surface from a known height above the cell lip, allowing the height of the overflowing froth to be calculated (as shown in Figure 3). The accuracy of the laser, as stated by the manufacturer, was ±1.5 cm, however this is dependent on lighting conditions, environmental conditions (e.g. dust) and the characteristics of the surface off which the laser pulse is reflected. Prior to the addition of the ultramarine dye particles, the penetration of the laser through the froth surface was clear; the dye particles inhibited this effect producing a more opaque surface. Additional checks using alternative methods were carried out that showed the accuracy was sufficient for the system.

For both the froth velocity and overflowing froth height measurements, filtering on the raw data was carried out within 5 s time periods, removing data points outside 1 standard deviation of the previous ten data points. Additional filtering was added between 5 s intervals to account for sudden deviations in the data; this was implemented by comparing the mean velocity or froth height from a given 5 s period with the previous 30 s, removing points outside 2 standard deviations and replacing with an average. A minimum standard deviation was set in order to avoid the system converging to a constant average value.

3. AIR RECOVERY CHARACTERISATION

Before implementing a control system, characterisation of the flotation cell was required to:

1) Establish the repeatability over the duration of a test and between tests
2) Determine the air recovery response as a function of air rate for modelling of the system and to determine where the control system should converge
3) Examine the dynamic response of froth velocity and overflowing froth height following a change in air rate for modelling of the system

Ten tests were carried out to investigate the variation in overflowing froth velocity and froth height over a range of 13 air rates from 130 to 190 lpm, equivalent to superficial gas velocity, J_g, of 1.10 cm/s to 1.61 cm/s respectively. The air rates were randomised for each test to avoid any system hysteresis affecting the results. The average froth velocity and overflowing froth height over all 10 tests are shown in Figure 4 and Figure 5 respectively. The error bars denote one standard deviation from the mean, with the dashed lines representing the 95% confidence interval.

The froth velocity is shown to increase with increasing cell aeration, with a plateau after 165 lpm. The velocity appears to drop at the very highest air rate, however this may be within experimental error. This is the inverse of the trend in overflowing froth height as cell aeration increases, which decreases sharply from 4.1 cm at 130 lpm to 3.6 cm at 160 lpm, before also remaining stable as the air continues to increase. The final increase in overflowing froth height at higher air rates (above 180 lpm) appears to be a real trend.

In industrial (i.e. particle stabilised) flotation systems, at low air flowrates, bubble surface area flux in the pulp phase is low relative to higher air rates (Hadler et al., 2010) resulting in an increased coating of hydrophobic particles. These more heavily-laden bubbles rise slowly in the froth phase due to the low superficial gas velocity, allowing drainage of entrained liquid and unattached particles. This slow-moving froth tends to have higher overflowing froth heights due to the increased bubble stability. As air rate increases, so too does the froth mobility; this gives rise to an increase in froth velocity but the increased mobility and higher liquid content yields a lower overflowing froth height. At higher air rates, the coating of stabilising particles on the bubbles decreases in the froth, and a fast flowing froth with unstable (bursting) bubbles is produced. This has been shown industrially (Hadler et al., 2006; Hadler and Cilliers 2009), however for the experimental system in this study, it must be recalled that it is two phase; the froth is not particle stabilised. It is important to note, however, that the froth behaviour for this laboratory system is similar to that observed industrially, suggesting that the reagent addition and viscosity modifier yield similar behaviour.

The shape of the trends in froth velocity and overflowing froth height with increasing air rate both appear to fit a quadratic curve, however it must be highlighted that the response of the froth to higher air rates than those tested here will not necessarily follow the same trend; as air rates increase further, the liquid content of the froth continues to increase until the region of the pulp-froth interface becomes indistinct and overflows the cell lip.

The froth velocity and overflowing height are combined to produce an air recovery curve. In order to clarify trends in air recovery with changing air rate, the average and standard deviation of the froth velocity and overflowing froth height are used to calculate an average air recovery and standard deviation per air rate. This is given in Figure 6, where it can be seen that a broad peak in air recovery occurs as the cell air rate increases. While the scatter in the data is large, the general trend can be observed, and this agrees with air recovery trends obtained industrially (Hadler and Cilliers, 2009). The peak in air recovery is the combined effect of increased froth mobility but decreased bubble stability as aeration increases. The magnitude of the measured air recoveries are somewhat higher than those measured at industrial sites (e.g. Smith et al., 2010), which are typically below 40% for rougher cells. It must be remembered, however, that the froth in the laboratory system is not particle stabilised and is operating in a high humidity environment. Moreover it is the trends that are important in the context of PAR and the measured peak in air recovery shows that this two phase froth exhibits trends similar to those measured industrially. This demonstrates its suitability as a system on which to test a PAR-based control system.

The large 95% confidence intervals are caused by the high standard deviation in the air recovery. Statistical analysis of the significance of the peak using Student's t-test shows that the air recovery at 160 lpm is statistically higher than at both 130 lpm and 190 lpm at the 95% confidence level. Consequently, while the peak in air recovery is statistically significant, the specific air rate at which PAR is obtained is less clear. In the following control system, an air rate in the region of 150 – 160 lpm is accepted as the PAR air rate, in light of the lack of a distinct PAR air rate.

4. DEVELOPMENT OF AN OPTIMISING FLOTATION CONTROL SYSTEM

Based on the PAR concept, a simplified flowsheet of the proposed control system is given in Figure 7, where the measurements for air recovery are taken, data filtering is applied and air recovery is calculated, before the controller establishes the next air rate. The PAR control system is based on direct search methods, which require no modelling of the system and do not rely on artificial intelligence. Specifically, the Generating Set Search (GSS) method was used to develop the PAR control system, and the algorithm used is outlined more fully in Figure 8.

In Figure 8, it is worth noting that the achievable step size sets both upper and lower limits (task 1), and for each step taken a check is performed to ensure the new value of the independent variable is within the constrained region (task 2). The method by Kolda et al. (2003) also recommends that for a specific step resulting in an increase in the objective function, the next step size should be increased to ensure the maximum is reached more quickly (true case of task 3). The simplicity and non-assuming nature of this optimisation method make it very versatile and robust, and it has been mathematically proven to converge on all continuous, differentiable, smooth functions. Kolda et al. (2003) further mentions that this method is able to handle noisy data, although there is no mention of the system being able to handle disturbances to the objective function.

It is important to note that only one variable, air flowrate, will be manipulated by the PAR control algorithm. Direct search methods are frequently used on systems where multiple parameters are altered, and for these cases the inter-dependencies of the manipulated variables need to be considered. For example, on a self-aerated flotation cell, changing the froth depth also results in a change in the air flowrate. Additionally, many direct search optimising methods are designed to converge to the maximum of a constant unknown function and then stop. Industrially, the PAR curve as a function of air flowrate is not constant and will change with varying feed, therefore for the algorithm tested here, the system continuously seeks to improve the air recovery even when the minimum step size has been reached, to accommodate for any changes in PAR or disturbances to the system.

The data from the baseline air recovery curve obtained during the characterisation phase of testing was used to build a dynamic model of the system on which the proposed optimising control algorithm was tested initially. The required information for the model included:

- The behaviour of v and h at steady state conditions.
- The behaviour of v and h under dynamic conditions.
- The amount of variance/noise expected in the v and h raw data.

Following the acquisition of all the required data, it was possible to model the behaviour of the froth phase at both steady state and dynamic conditions for a given change in Q_a. The addition of noise to the predicted trend was achieved using an inverse Gaussian distribution function, and a series of simulated random variables to dictate how much deviation from the mean should be applied to each data point. This allowed for the predicted results to be simulated with the typical standard deviation that is present in the experimental data, which for the laboratory system was 0.53 cm/s and 0.14 cm for v and h respectively. Note that these values are the standard deviation from the mean for any given test at a given air rate. Importantly, the modelled noise is randomly generated, which means that no two tests were the same and the degree to which the simulated data deviated from the expected trend varied between tests (as was the case on the laboratory system).

Using the dynamic and steady state models for v and h as a function of Q_a with and without noise enabled the simulation of the GSS method. The air flowrate settings used in the simulation trials were as follows:

- Q_a range allowed: 130 – 190 lpm
- Maximum Q_a step size: 25 lpm

8

- Minimum Q_a step size: 5 lpm
- Initial Q_a step size: 20 lpm (arbitrarily chosen as 33% of full Q_a range)

The allocated PAR control timing regime was as follows:

- Stabilisation time: 60 min (allowing h decay to become linear)
- Measurement time: 9 min (time allocated for v and h measurement)
- Wait time: 6 min (time allowed to reach steady state after Q_a change, based on kinetics of change)
- Total run time: 6 hours (i.e. 1 hour stabilisation time, 5 hours of the system being optimised)
- Step size expansion factor: 120%
- Step size contraction factor: 50%

The results of three simulated PAR control runs (with noise added to simulated data), all starting from an air flowrate of 130 lpm, therefore lower than the PAR air rate, are given in Figure 9, where the step change in air flowrate and air recovery response are shown. In each case, it can be seen that the air rate is changed such that the air recovery increases towards PAR. The difference in the three simulations are the result of noise added to the data; it can be seen that the top two simulations yield the same response in the initial 4 air rate changes, however the third simulation shows a drop to 130 lpm and a corresponding drop in air recovery. Similar simulations were carried out for cases starting at an air rate above PAR (190 lpm) and at the PAR air rate (160 lpm).

A summary of the convergence times are given in Table 1 and it can be seen that the shortest time to obtain PAR was 105 minutes, or 7 iterations, which was achieved when starting at air rates below the PAR air rate. This highlights the slow approach of the GSS direct search method, however in all cases, PAR was obtained, taking into account experimental noise. There is clear room for improvement on this technique, however it is robust in finding the PAR air rate.

5. CONTROL SYSTEM IMPLEMENTATION AND TESTING

The GSS model tested on the simulated data was applied to the experimental system using LabVIEW. Three cases were tested; starting at an air rate below the PAR air rate, at an air rate higher than the PAR air rate, and starting at the PAR air rate with an artificial disturbance introduced into the system. Once the PAR control system was started, no interaction was made with the system at all, with the exception of the case where the effects of a disruption to the objective function were being investigated.

The first test of the PAR control system on the laboratory flotation cell was for the case where the initial air rate is lower than the expected PAR air rate. The results of the trial are shown in Figure 10, where the numbered points show the air rates and corresponding measured air recoveries in the sequence set by the control system. The air recoveries obtained are compared to the results from the PAR characterisation studies, showing good agreement between the values obtained during the PAR control test and the averaged values from the preliminary tests.

From Figure 10, it can be seen that the GSS-based PAR control system was able to rapidly drive the laboratory system towards PAR (points 1 to 3), and then maintain the system near PAR. The majority of the results lie between 150 – 170 lpm, this despite some considerable noise in the measured data (examples being points 4 and 7). Moreover, these results are in good agreement with the simulated results, although the observed noise was greater in the experimental results. The greater amount of noise in the experimental data compared to the simulated data is likely to have originated in the standard deviation in the simulations being an averaged standard deviation. Secondly, the model is not sensitive to all of the variables that affect the laboratory system, leading to unexpected disturbances in the experimental data. For example, the humidity in the enclosure tended to rise slowly over time, especially when the cell operated at higher air rates, yielding slight increases in the measured air recoveries. This observation agrees with the study of Li et al. (2010).

In a second trial of the control system, the initial air rate was set above the PAR air rate. The results of this trial can be seen in Figure 11. It is interesting to note that during this trial, the flotation system was not stable across the 6 hour test duration, as demonstrated by the increasing air recovery over time. This originated in the overflowing froth velocity increasing over the test, suggesting that the frother addition rate was too high for this experiment. This highlights the sensitivity of flotation froths, both for this system and industrially, to small changes in one of the many variables that effect froth stability.

Nevertheless, despite this instability in the laboratory system, the GSS-based PAR control system was still clearly able to drive the system towards PAR. This example further illustrates the robustness of the GSS algorithm and its ability to continue to optimise the system even when the objective function is slowly changing with time, as will often be the case for industrial flotation systems where changes in feed properties occur frequently.

From the results of the simulations on the performance of the control system, the GSS algorithm was found to handle efficiently (albeit slowly) both starting at the PAR air rate and step change disturbances in the objective function. From the previous control trials, however, the artificially-added noise in the simulations is an underestimate of the experimental noise. In order to test fully the robustness of the GSS-based PAR control system, a trial was carried out starting at PAR and experiencing a large step change in the objective function, introduced by multiplying the overflowing froth height, h, by an arbitrarily selected factor in order to drive away the air rates from the PAR position. The results are shown in Figure 12.

From Figure 12, it can be seen that steps 4, 5 and 6 were the points in which the manual disturbance was implemented. Following the disturbance, the control system was able to drive the system back to PAR, converging between points 10-14. This particular experiment was repeated three times, in each case the GSS-based PAR control system performed similarly, despite the level of experimental noise. This is further testament to the simple, robust strategy employed by the GSS direct search method; that is, to only consider the last two data points, and constantly direct the air flowrate in the direction that results in the air recovery increasing. Inclusion of AI into a PAR control system may well result in increased efficiency.

These results demonstrate, for the first time, a flotation optimisation control system based on a direct search approach. It should be noted that this has been implemented for a single cell only, where industrially, cells operate in series with the performance of the first cell affecting the feed into the second. The advantage of this approach and the PAR concept as a control strategy is that the optimal air rate for given feed and reagent conditions will be targeted always.

6. CONCLUSIONS

Optimising control of froth flotation presents a challenge for industrial operations, not least due to instabilities in industrial circuits arising from variations in feed material. Air recovery is the fraction of air entering a flotation cell that overflows as unburst bubbles and is used to quantify froth stability. Air recovery is known to pass through a peak as cell aeration is increased, and it has been shown that operating flotation cells at the air rate that yields this peak air recovery (PAR) results in improved separation performance, whether through high recovery, higher grade, or both.

The concept of PAR introduces a target, or objective function, for use in the control of flotation cells. In this paper, a large scale (70 l) continuous laboratory flotation system was designed with a reagent system that demonstrated characteristics similar to those found industrially, that is, exhibiting a peak in air recovery with increasing air rate. Following characterisation tests to establish an average air recovery curve, the response of the system to changes in air rate, both under dynamic (i.e. changing air rate) and steady state conditions was modelled to allow a suitable platform on which to test potential control algorithms. The Generating Set Search (GSS) direct

search method was shown to be a slow but robust method, capable of finding the PAR air rate using the simulated data when starting below, above and at PAR.

The GSS control system was trialled on the laboratory rig, showing in each case that PAR could be obtained and maintained, even under conditions in which air recovery was varying over time. This demonstrates that suitability of such a control system to the changing conditions observed at most industrial flotation operations.

This study has demonstrated for the first time a novel optimising control strategy and system based on PAR for flotation cells.

ACKNOWLEDGEMENTS

REFERENCES

Aldrich C., Moolman D.W., Gouws F.,S., Schmitz G.P.J., (1997). Machine learning strategies for control of flotation plants. *Control Engineering Practice*. 5(2), 263,269

Aldrich C., Marais B.J., Shean B.J., Cilliers J.J., (2010). Online monitoring and control of froth flotation systems with machine vision: A review. *International Journal of Mineral Processing*, 96 **(1-4)**, 1-13

Bergh L.G., Yianatos J.B., Acuna C.A., Perez H., Lopez F., (1999). Supervisory control at Salvador flotation columns. *Minerals Engineering*, 12 (7), 733-744

Bergh L.G., Yianatos J.B., (2011). The long way toward multivariate predictive control of flotation processes. *Journal of Process Control,* **21**, 226–234.

Brito-Parada P.R., Cilliers J.J., (2012). Experimental and numerical studies of launder configurations in a two-phase flotation system. *Minerals Engineering*, **36-38**, 119-125

Dey A., Kapur P.C., Mehrotra S.P., (1989). A search strategy for optimization of flotation circuits. *International Journal of Mineral Processing*, **26** (1-2), 73-93

Hadler K., Barbian N., Cilliers J.J., (2006). The relationship between froth stability and flotation performance down a bank of cells. *In: Proceedings of the XXIII IMPC, Istanbul, Turkey*

Hadler K., Cilliers J.J., (2009). The relationship between the peak in air recovery and flotation bank performance. *Minerals Engineering*, **22**, 451–455.

Hadler K., Smith C.D., Cilliers J.J., (2010). Recovery vs. mass pull: The link to air recovery. *Minerals Engineering,* **23**, 994–1002.

Hadler K., Greyling M., Plint N., Cilliers J.J., (2012). The effect of froth depth on air recovery and flotation performance. *Minerals Engineering*, **36-38**, 248-253

Hodouin D., Jamsa-Jounela S.-L., Carvalho M.T., Bergh L., (2001). State of the art and challenges in mineral processing control. *Control Engineering Practice*, **9** (9), 995-1005

Jahedsaravani A., Marhaban M.H., Massinaei M., Saripan M.I., Noor S.B.M., (2016). Froth-based modelling and control of a batch flotation process. *International Journal of Mineral Processing*, **146**, 90-96

Karelovic P., Putz E., Cipriano A., (2015). A framework for hybrid model predictive control in mineral processing. *Control Engineering Practice*, **40**, 1-12

Kaartinen J., Hätönen J., Hyötyniemi H., Miettunen J., (2006). Machine-vision-based control of zinc flotation – A case study. *Control Engineering Practice*, 14(12), 1455-1466

Kolda T.G., Lewis R.M., Torczon V., (2003). Optimisation by direct search: new perspectives on some classical and modern methods. *Society for industrial and applied mathematics*. **V45**, 3, pp. 385-482.

Laurila H., Karesvuori J., Tiili O., (2002). Strategies for instrumentation and control of flotation circuits. *Mineral Processing Plant Design, Practise and Control*, **Volume 1**, 2174-2195

Li X., Shaw R., Stevenson P., (2010). Effect of humidity on dynamic foam stability. *International Journal of Mineral Processing*, **94** (1-2), 14-19.

Maldonado M., Desbiens A., del Villar R., (2009). Potential use of model predictive control for optimizing the column flotation process. *International Journal of Mineral Processing*, **93** (1), 26-33.

Marlin, T. E., (2000). Process control: designing processes and control systems for dynamic performance. *McGraw-Hill.*

Mathe Z.T., Harris M.C., O'Connor C.T., Franzidis J.-P., (1998). Review of froth modelling in steady state flotation systems. *Minerals Engineering*, **11** (5), 397–421.

McKee D.J., (1991). Automatic flotation control — a review of 20 years of effort. *Minerals Engineering* **4** (7–11), 653–666

Neethling S.J., Cilliers J.J., (2003). Modelling flotation froths. *International Journal of Mineral Processing*, **72** (1-4), 267-287

Neethling S.J., Cilliers J.J., (2008). Predicting air recovery in flotation cells. *Minerals Engineering*, **21**(12-14), 937-943

Norori-McCormac A., Hadler K., Cilliers J.J., (2014). Peak Air Recovery: An investigation into the effect of particle size. *In: Proceedings of the XXVII IMPC, Santiago Chile*

Nunez F., Tapia L., Cipriano A., (2010). Hierarchical hybrid fuzzy strategy for column flotation control. *Minerals Enigneering*, 23(2), 117-124

Osorio D., Pérez-Correa J.R., Cipriano A., (1999). Assessment of expert fuzzy controllers for conventional flotation plants. *Minerals Engineering* **12** (11), 1327–1338

Pugh R.J., (1996). Foaming, foam films, antifoaming and defoaming. *Advances in Colloid and Interface Science*, **64**, 67-142

Putz E., Cipriano A., (2015). Hybrid model predictive control for flotation plants. *Minerals Engineering*, **70**, 26-35

Qu X., Wang L., Nguyen A.V., (2013). Correlation of air recovery with froth stability and separation efficiency in coal flotation. *Minerals Engineering,* **41**, 25-30

Shean B.J., Cilliers J.J., (2011). A review of froth flotation control. *International Journal of Mineral Processing*, **100** (3-4), 57-71

Smith C.D., Hadler K., Cilliers J.J., (2010). Flotation bank air addition and distribution for optimal performance. *Minerals Engineering*, **23**, 1023-1029

Supomo A., Yap E., Zheng X., Banini G., Mosher J., Partanen A., (2008). PT Freeport Indonesia's mass-pull control strategy for rougher flotation. *Minerals Engineering,* **21**, 808–816.

Wang L., Peng Y., Runge K., Bradshaw D., (2015). A review of entrainment: Mechanisms, contributing factors and modelling in flotation. *Minerals Engineering*, **70**, 77-91

Woodburn E.T., Austin L.G., Stockton J.B., (1994). Froth based flotation kinetic model. *Chemical Engineering Research and Design*, **72** (A2), 211-226

Xiong Q., Jutan A., (2003). Continuous optimization using a dynamic simplex method. *Chemical Engineering Science*, **58**, 3817-3828

Zhu J., Gui W., Yang C., Xu H., Wang X., (2014). Probability density function of bubble size based reagent dosage predictive control for copper roughing flotation. *Control Engineering Practice*, **29**, 1-12
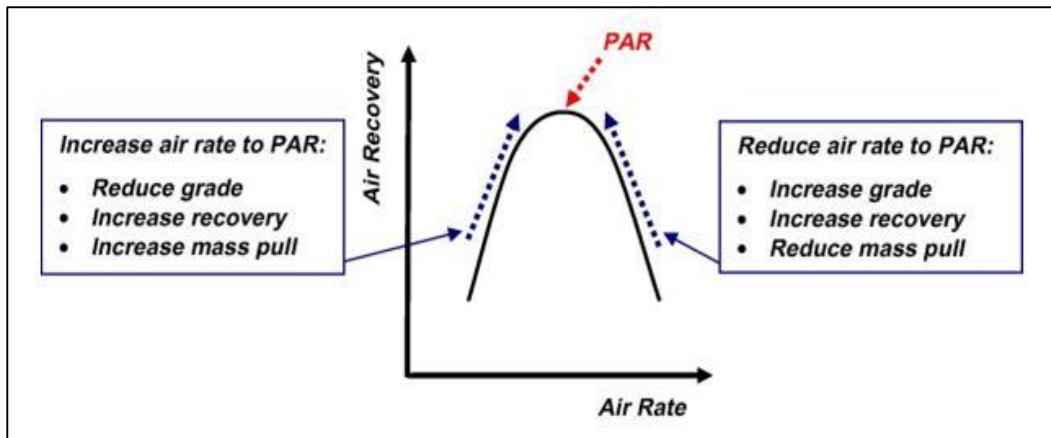
Figure 1: Effect of air rate on flotation performance with respect to Peak Air Recovery (PAR) (Hadler et al., 2010)
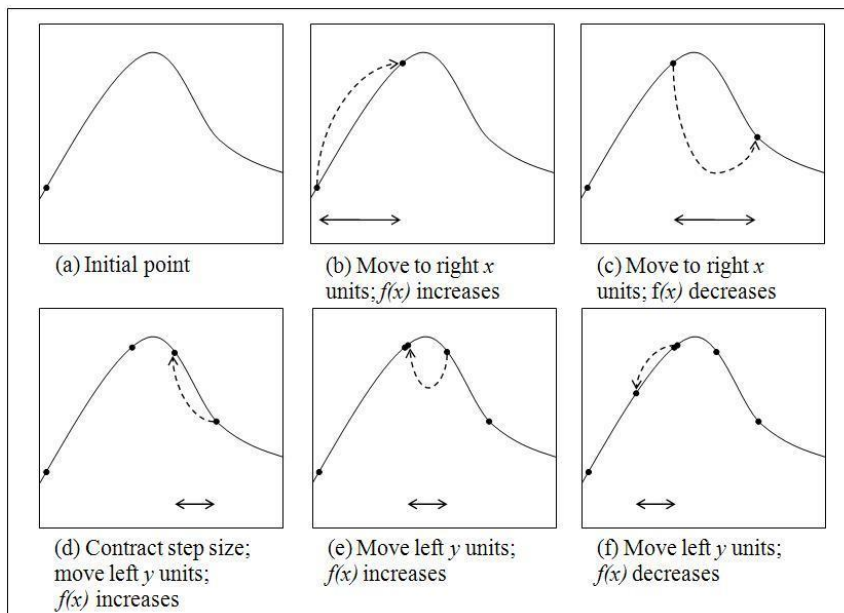
Figure 2: Simplified example of the GSS method

Figure 3: Diagram of the custom-built laboratory flotation cell, including overflowing froth height measurement by laser and camera for image analysis of froth surface
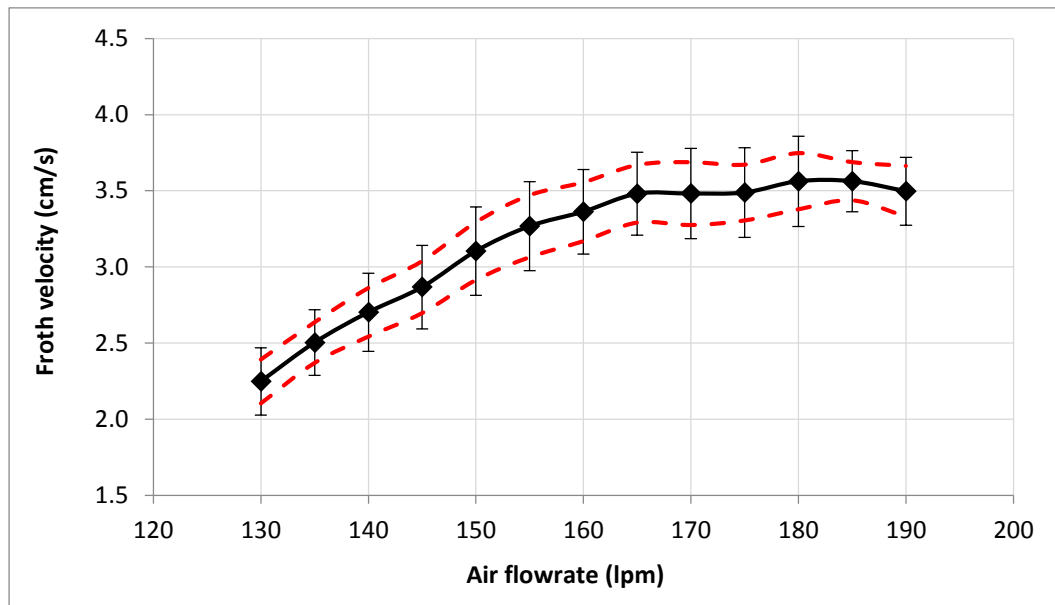

Figure 4: Average overflowing froth velocity as a function of air rate over 10 repeats. Vertical error bars represent one standard deviation in the data (across different tests), and the dashed lines represent the respective upper and lower 95% confidence limits.
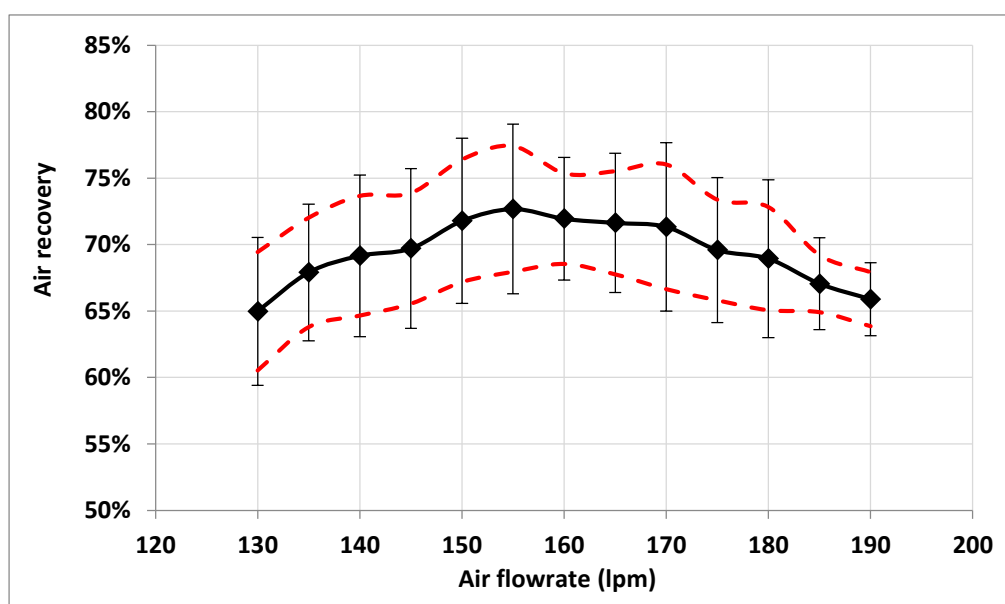
Figure 5: Average overflowing froth height as a function of air flowrate over 10 tests. The vertical error bars represent one standard deviation in the data (across different tests), and the dashed lines represent the respective upper and lower 95% confidence limits.

Figure 6: Average calculated air recovery values as a function of air rate over 10 tests. The vertical error bars represent one standard deviation in the data, and the dashed lines represent the respective upper and lower 95% confidence limits.
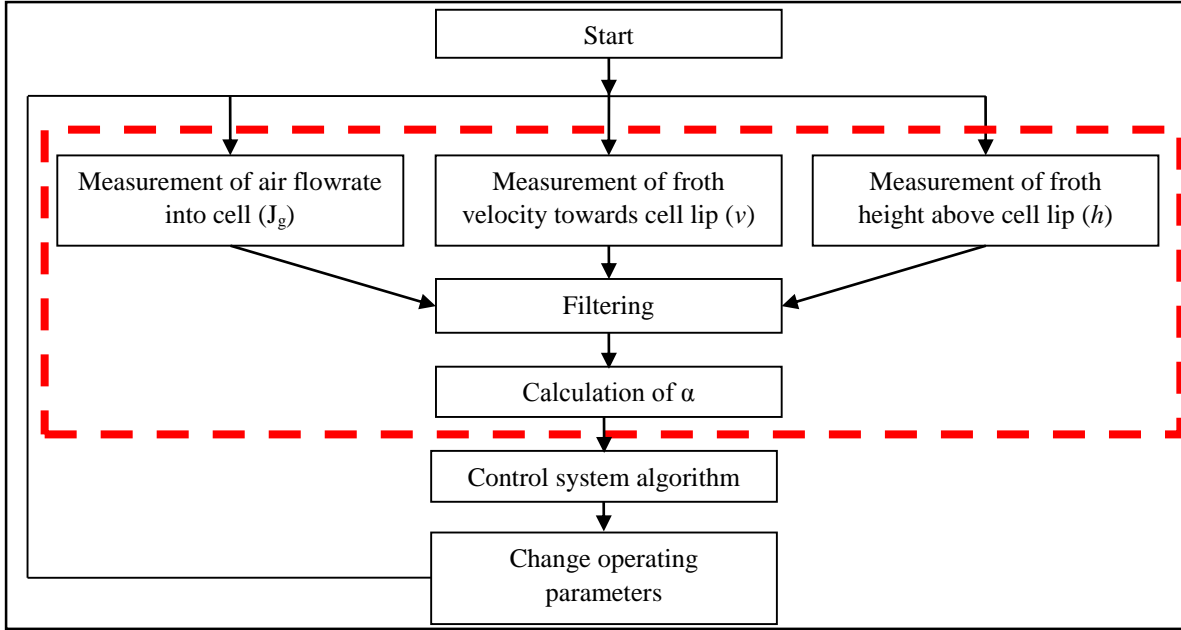
Figure 7: Simplified block diagram of PAR control system; the dashed box highlights the air recovery measurement module
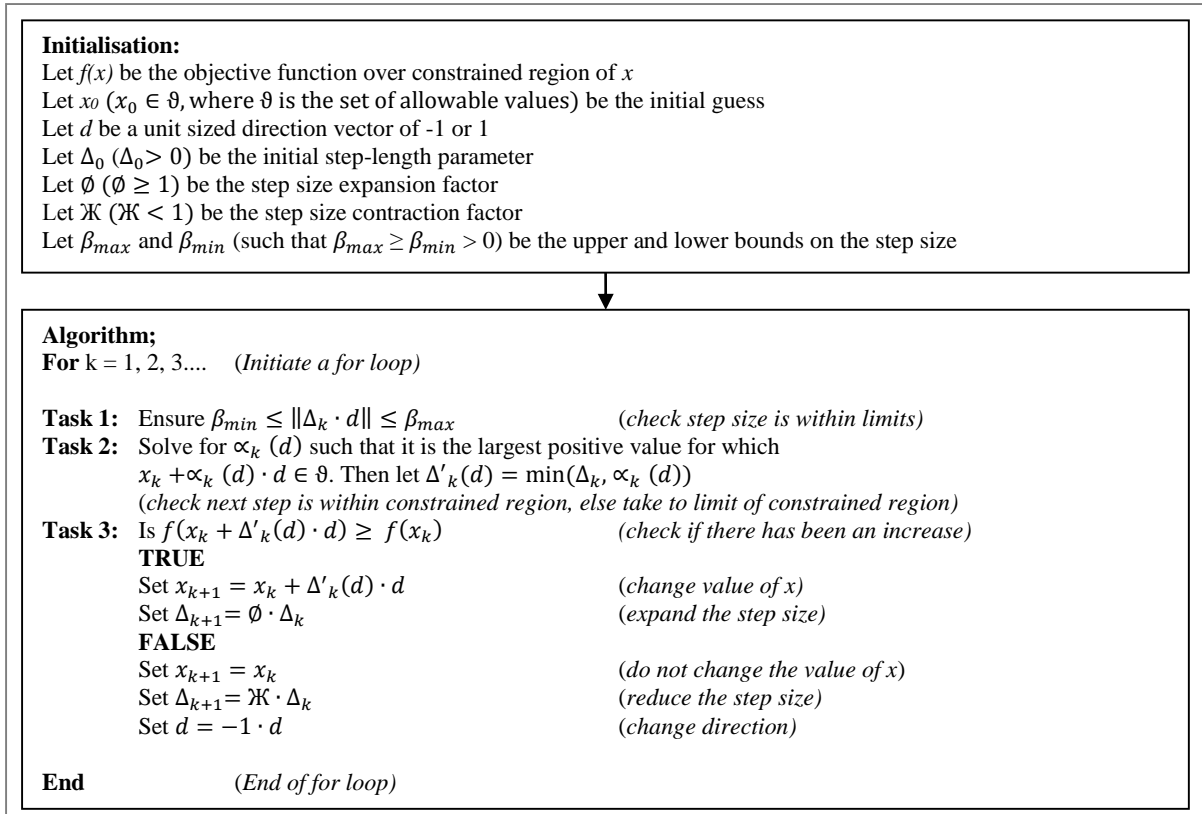

Figure 8: Modified generalised constrained region GSS algorithm from Kolda et al. (2003)
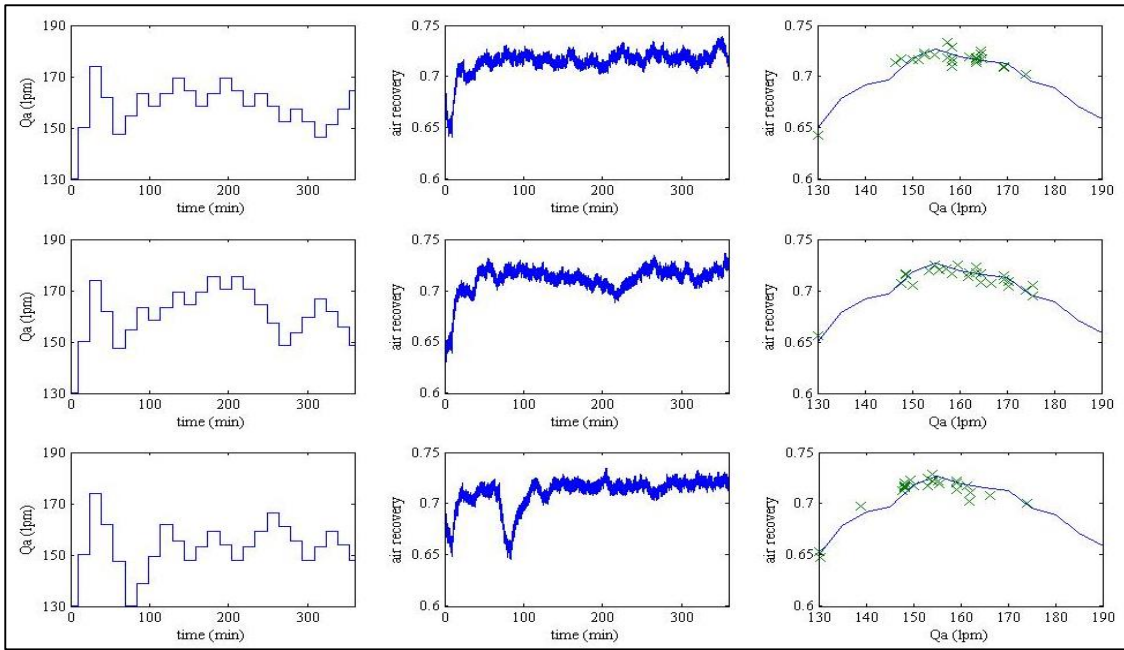
Figure 9: Results of GSS simulations. Q_a initially set at 130 lpm (lower than PAR air rate). Each row shows Q_a (lpm) as a function of time (min) (left hand column) (centre column), air recovery as a function of time (min), and air recovery as a function of Q_a (lpm)(right hand column).


Figure 10: Expected air recovery as a function of air rate (heavy dashed line) and air recoveries obtained using the GSS-based PAR seeking control system starting from an air rate lower than the PAR air rate. The order in which the air rates were selected by the control system is given by the numbered data points
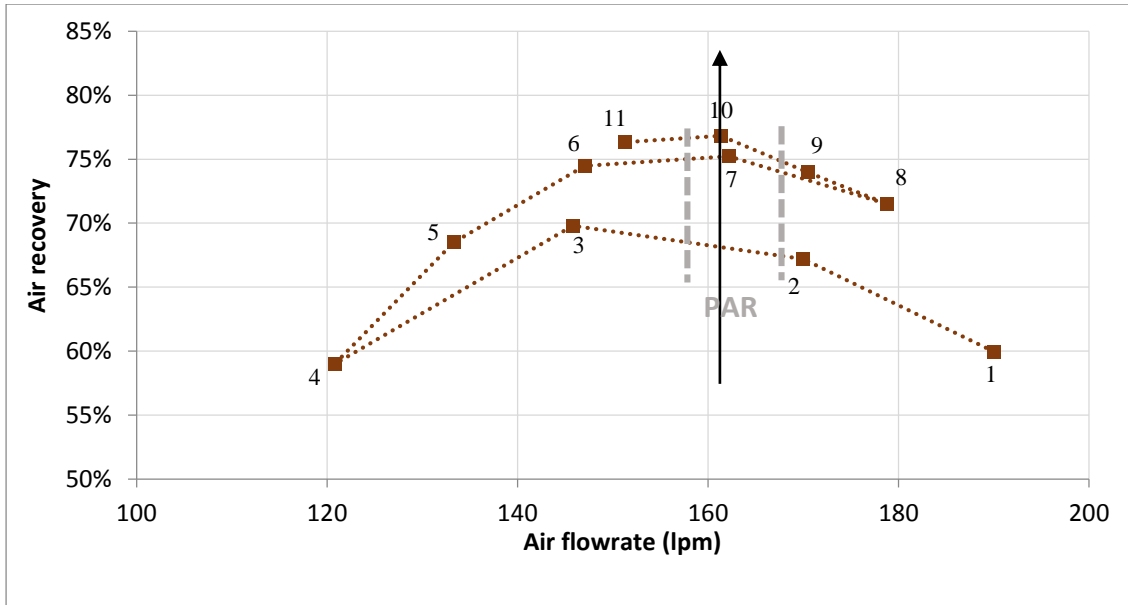
18

Figure 11: Expected air recovery as a function of air rate (heavy dashed line) and air recoveries obtained using the GSS-based PAR seeking control system starting from an air rate higher than the PAR air rate. The vertical arrow indicates the direction in which air recovery increased over the experimental duration.
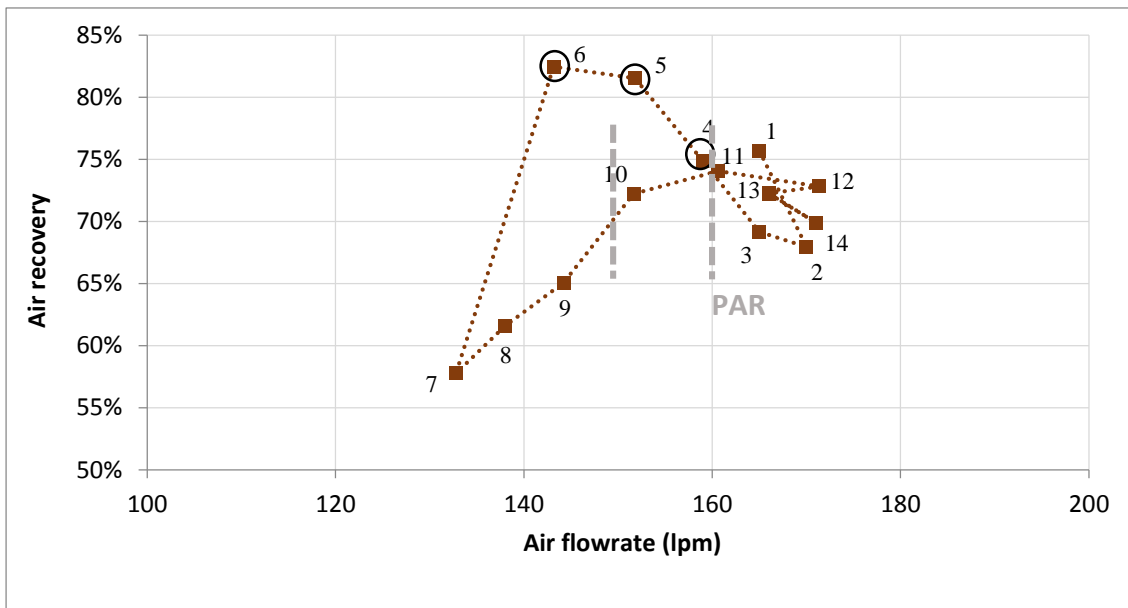

Figure 12: Expected air recovery as a function of air rate (heavy dashed line) and air recoveries obtained using the GSS-based PAR seeking control system. The circled data points are those manually changed to induce a system disturbance by moving the air recovery away from PAR.

Table 1: Time to converge at PAR for the simulation with and without noise starting from below, at and above PAR

	Time to reach PAR (without noise)	Number of iterations	Time to reach PAR (with noise) (three repeats at each condition)
Time for convergence (min) starting below PAR	105	7	First convergence times at 105, 105, 195 min
Time for convergence (min) starting above PAR	225	15	First convergence times at 105, 120, 195 min
Time for convergence (min) starting at PAR	180	12	First convergence times at 165, 150, 135 min