Tank design modifications for the improved performance of froth flotation equipment

A. J. Morrison, P.R. Brito-Parada, J.J. Cilliers

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

Abstract

Turbulence plays a complicated role in froth flotation: It is desirable because it helps to promote bubble-particle interactions in the pulp, and to increase the specific surface area of the air bubbles introduced into the tank, but it is undesirable because it also promotes the detachment of valuable mineral particles from bubbles, and exacerbates entrainment-inducing mixing at the pulp-froth interface.

In this work, an experimental program was undertaken to optimise the performance of an easily-retrofittable froth phase insert designed to isolate the froth, and the pulp-froth interface, from the effects of pulp-phase turbulence in a laboratory-scale froth flotation tank. The outcome of this program is a range of optimal insert designs that result in up to a 19% improvement in recovery from the tank without reducing the grade of the concentrate, or up to 3% improvement in the grade without reducing recovery from the tank.

Keywords: froth flotation, froth stability, grade-recovery curve

1. Introduction

Froth flotation is a separation technique that exploits the differences in surface hydrophobicity between a valuable mineral to be recovered and the various waste materials– called gangue– from which it is to be isolated (Wills, 2006). A hydrophobic surface, when it comes into contact with the interface between a hydrous solution and a gas will tend to align itself so as to minimise its exposure to water (Gupta and Yan, 2006). Thus, in an aerated pulp, sufficiently-small particles of a hydrophobic mineral will– under the right conditions– tend to attach themselves to the thin films surrounding any bubbles that they chance to

*Corresponding author
Email address: ajm612@ic.ac.uk (A. J. Morrison)

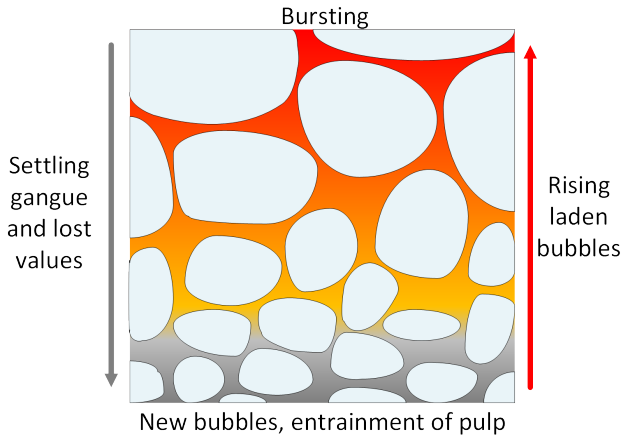


Figure 1: The structure of a flotation froth.

encounter . Since bubbles would naturally be rising through such a slurry, and denser-than-water solid particles falling, this provides a mechanism by which to segregate hydrophobic mineral particles to the top, while allowing the less hydrophobic gangue particles to settle to the bottom of the aerated tank.

Particle-laden bubbles that reach the surface of the pulp join a three-phase froth layer consisting of solids, liquid and gas . As bubbles enter the froth layer, they non-selectively trap pulp between themselves in a process called entrainment . This entrained pulp, while richer in valuable mineral than pulp deeper in the tank due to the selection effect of rising laden bubbles, still contains gangue particles. Thus, froth flotation is incapable of perfect selection of hydrophobic mineral particles from the pulp.

However, as more and more bubbles enter the froth, older bubbles are pushed higher up in the froth layer. As they are, the entrained pulp trapped between them begins to drain out under the effect of gravity, and the froth becomes drier. Drier bubbles in the froth coalesce into larger ones, reducing their combined surface area to volume ratio. This both reduces their carrying capacity for attached solids, and the carrying capacity of the froth for entrained pulp (Wills, 2006). This phenomenon, called froth coarsening , is shown schematically in figure 1.

Eventually, the coalescing rate of the bubbles in the froth is so high that the froth cannot grow any more, and a froth surface forms. Bubbles that reach the surface of the froth dry out even more quickly and burst, shedding their attached particles in the process. In order to recover the concentrated valuable material from the froth, it is necessary to collect the froth before its bubbles burst, generally by allowing it to overflow the tank.

Since the gangue is only in the entrained pulp, and not attached to the bubbles,

the deeper the froth is, the higher the ratio of valuable mineral to gangue in the overflowing concentrate . The proportion of valuable minerals in the overflowing froth is called the concentrate grade. However, the deeper a froth is, the lower the surface area to volume ratio of its overflowing bubbles, and thus their carrying capacity for attached valuable minerals per unit volume of air supplied to the tank. This reduced carrying capacity reduces the total amount of valuable mineral that is eventually recovered from the pulp. This concentrate grade and recovery are the main measures of the performance of a froth flotation system (Wills, 2006; Napier-Munn, 2014).

The maximum height that a froth can achieve is related to the stability of the froth (Gourram-Badri et al., 1997), which is influenced by factors such as the size and surface hydrophobicity of particles in the pulp , the rheology of the entrained liquid in the froth , and the influence of the hydrodynamic environment in the pulp . In particular, (Morrison, 2017), found that turbulence at the pulp-froth interface had a significant effect on stability of the froth, as measured by a range of proxies such as froth height and overflowing volume, and the specific air recovery from the system . In this paper, a design modification will be introduced, whose aim is to improve the metallurgical performance of a mechanical froth flotation tank by reducing the effect of turbulence on the pulp-froth interface.

2. Mechanical froth flotation

At its simplest, a mechanical froth flotation machine consists of a mechanically-agitated tank into which air is either introduced through a sparger, or inducted from the surface through the action of the agitator. A pulp of mixed valuable material and gangue enters the tank through one or more inlet ports, and is removed from the bottom of the tank as tails, and from the top of the tank as a froth. Through these actions, the froth has a higher grade of valuable minerals, and the tails has a lower one. The froth that overflows the lip of the flotation tank is collected in a more or less complicated system of troughs, called launders, and redirected for either further concentration or refining.

The agitator, or impeller, in a mechanical flotation tank is what differentiates it from a flotation column, in which mixing only occurs through the action of rising bubbles . A mechanical impeller fulfils two important duties: It keeps the solids in the pulp suspended long enough for the valuable mineral particles to interact with and attach to bubbles, and it creates a turbulent region which shears and disperses the bubbles introduced below it, increasing their surface area, and thus the likelihood of bubble-particle interactions .

However, a mechanical impeller also has two main disadvantages in froth flotation. The first is that, in order to efficiently suspend particles, it must also create mixing loops in the pulp, which carry gangue from the inlet to the pulp-froth interface, and bubbles from the interface back into the pulp . This reduces

the separation effect that rising, laden bubbles would otherwise have, and thus increases the importance of the froth in improving the grade of the concentrate. The need to counteract impeller-induced transaxial and axial mixing loops is the rationale for varying industrial impeller design, and for modifying the shape and internal structure of froth flotation tanks.

The second disadvantage of an impeller in a froth flotation tank is its effect on the pulp-froth interface. Both the mixing loops used to suspend particles, and the shear induced by the impeller to disperse and size-reduce bubbles, lead to a turbulent environment that can have a deleterious effect on the froth stability by destabilising the interface between the pulp and the froth . This reduction in froth stability leads to a greater coalescence rate of bubbles, and thus to a greater draining rate of entrained pulp, reducing the amount of solids recovered from the pulp as overflowing froth. A turbulent pulp-froth interface can also increase the mixing of the pulp and the froth, and thus reduce the grade of the final concentrate. This process is called churn . Confining impeller-induced turbulence to where it is necessary has led to various design improvements such as the introduction of rotor-stator impellers, and flow-modifying draft tubes and impeller cages.

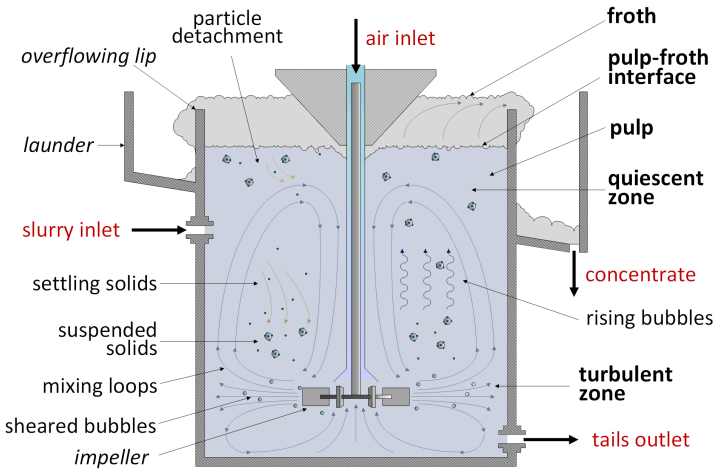


Figure 2: The hydrodynamic environment inside a flotation tank, showing the behaviour of bubbles and particles in the turbulent, quiescent and froth zones.

Figure 2 is a schematic of a typical mechanical flotation tank, indicating the turbulent pulp zone in which the effect of the impeller is strongest, and the quiescent zone in which most of the impeller-induced turbulence has dissipated. Figure 2 also shows the mixing loops that would be induced by an unmodified Rushton impeller .

An almost ubiquitous design feature of commercial flotation tanks is the enclosure of the impeller to create a concentrated turbulent zone [Mavros, 1992].

These turbulent zone enclosures are varied– from the rotor-stator design of, for instance, Outotec tanks , to the air disperser design of Wemco tanks – but all share the same rationale: that turbulence is desirable to promote bubble-particle interactions (Ralston et al., 1999), and to produce smaller bubbles [Napier-Munn,2014], but also undesirable because it promotes the detachment of valuable mineral particles from bubbles, and exacerbates entrainment-inducing churn at the pulp-froth interface (Jameson, 2012).

Clearly, whether or not turbulence is desirable varies from place to place in the tank. By enclosing the rotor in a stator, or the impeller assembly in an air dispersion cage, not only can the effect of the turbulent zone on the quiescent zone be more tightly restricted, but the interaction of the fast moving pulp with immovable structures also increases the shear in the turbulent zone, which improves the efficacy of the impeller in producing small bubbles. However, the competing requirement that the impeller keep solids in the pulp suspended long enough for them to come into contact with rising bubbles means that there is a limit to how effectively the turbulent region can be contained.

One possible design modification that would further help to segregate the turbulent and quiescent zones, but which has not yet been applied to mechanical flotation tanks, is the horizontal baffle. Horizontal baffles were introduced for column flotation by Kawatra and Eisele (1994). Column flotation is froth flotation in an impeller-less tank, uniformly sparged at its base, and relying on the buoyancy of the resultant bubbles to both separate the valuable mineral and gangue and to set up a mixing loop in the tank.

The strength of the mixing cells in flotation columns pulp rising at the centre of the column, dragged up by buoyant bubbles, and falling commensurably at the walls is determined by the size of the bubbles in the tank, with larger bubbles leading to more developed mixing. Strong mixing cells, as well as the momentum transferred by the arrival of large bubbles, result in a turbulent zone at the pulp-froth interface (Prince and Blanch, 1990). The resultant churn leads to entrainment of gangue, detachment of valuable minerals at the base of the froth.

Kawatra and Eiseles horizontal baffles consist of grids or perforated plates placed transaxially along the height of a flotation column in order to segregate it into many subregions. The size of the apertures in these horizontal baffles are such that fluid and bubbles of up to a certain size can pass through them, but that they disrupt the bulk mixing cells that would otherwise arise in a flotation column. Eisele et al. (1999) showed significant improvement in the concentrate grade produced by coal desulphurisation columns retrofitted with horizontal baffles. They attributed this to increased residence time and reduced churn in the column. There was no difference in coal recovery because desulphurisation is an upgrading process, and so the concentrate reports to the tails in a process called reverse flotation .

In a precursor to the present study, Norori-McCormac et al. (2017) found that

the insertion of a horizontal baffle at the pulp-froth interface of a 4-litre, bench-scale, mechanical flotation tank with a single-species pulp led to an increased indicative recovery. This result suggests that similar gains could be achieved at a more realistic scale.

3. Froth flotation performance

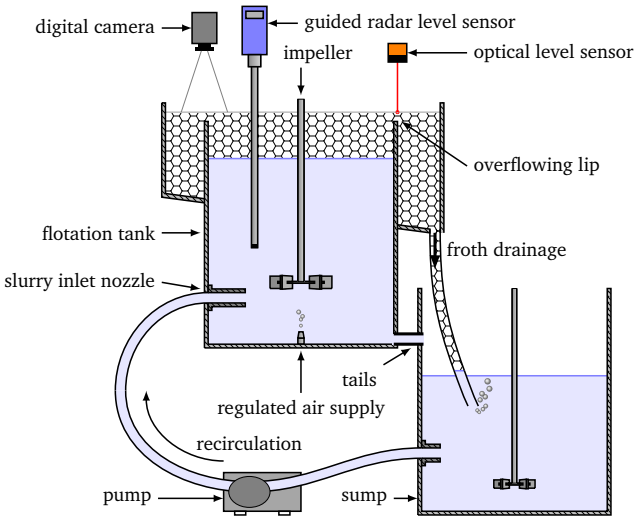


Figure 3: A schematic of the continuously-operated, laboratory-scale froth flotation circuit.

A schematic of the laboratory-scale froth flotation circuit used for this work is shown in figure 3. The circuit is built around a baffled, 87-litre cylindrical mixing tank that has been modified for froth flotation by the addition of a launder, sparger, and slurry inlet and outlet ports. The tank is agitated by a standard Rushton impeller mounted a third of the way up between its base and overflowing lip. Both the concentrate that overflows the laboratory-scale tank into its launder, and the tails that are drained from its base, flow into a sump to be remixed and pumped back into the tank. The result is a continuously-operable, laboratory-scale mechanical flotation circuit, capable of running stably and repeatably for 9 hours (Morrison, 2017).

The airflow rate into the tank, the impeller rotation rate, and the tails flow rate are all remotely controlled from a computer control system. The same control system is used to log the depth of the pulp-froth interface below the lip of the tank, using a guided radar level sensor (GRL); the height of the overflowing froth above the lip of the tank, using an optical level sensor (OLS); and the overflowing rate of the froth, using image analysis of the surface of the flowing froth. Details of the control and instrumentation fitted to the circuit, as well as the baseline flotation performance of the circuit, are given by Morrison (2017).

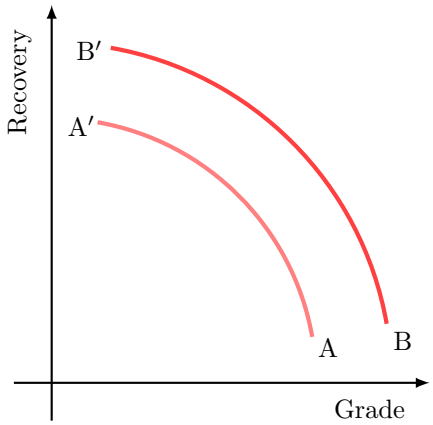


Figure 4: A pair of typical grade-recovery curves.

3.1. The metallurgical performance of a froth flotation tank

In mineral processing, it is not only important to achieve as pure a product as possible, but also to recover as much of the valuable mineral from the pulp as possible. If the system were completely non-selective, then all of the valuable mineral would be recovered in the concentrate, which would have the same grade as the raw feed. If it were perfectly selective, it would reject even ore particles containing a very small amount of gangue (Wills, 2006), resulting in a very pure concentrate, but also a lot of valuable mineral in the waste. Thus there is a tension between the selectivity and the recovery of the system.

The proportion of the total available valuable mineral that reports to the concentrate is called the recovery of the system [Wills, 2006]. Good recovery relies on removing the valuable mineral from the pulp, and into the froth, before it has an opportunity to be removed from the base of the tank as waste. This means increasing the probability of bubble-particle attachment in the pulp, removing these laden bubbles from the pulp into the froth before they have a chance to burst, and removing the enriched froth from the tank before it coalesces and begins draining valuable minerals .

Increasing the probability of attachment means either increasing the bubble surface area in the pulp, or keeping mineral particles suspended longer through increased mixing of the pulp. However, both of those increase the turbulence in the pulp in ways which decrease the grade produced by the froth. Removing laden bubbles before they burst, and before the entrained pulp has a chance to drain also reduces the grade that can be achieved by the froth, as was described above, and so optimising the performance of the froth flotation process is a balance between the purity of the concentrate its grade and the total amount of valuable mineral recovered (Napier-Munn, 2014).

This trade off is generally presented in terms of a grade-recovery curve, of the

form shown in figure 4. Optimising the performance of a froth flotation system then, can take one of two forms. Either a combination of operating conditions can be chosen so as to prioritise grade over recovery or vice versa, or a modification can be made to the system to simultaneously increase the possible combinations of achievable grade and recovery. This first procedure is the equivalent of moving along the grade-recovery curve between position A, where the grade is maximised, and position A′, where recovery is maximised. The second procedure is the equivalent of moving from grade-recovery curve A to curve B.

3.2. Single species proxies for grade and recovery

In order to minimise the effects of physical and chemical feed variability and ageing on the performance of the laboratory-scale tank , this work was carried out on a pulp containing only deionised water and hydrophobised silica particles. Since this pulp contained only valuable mineral and no gangue analogues, it is a single-species flotation pulp . Such a pulp is often used for initial reagent dosing and floatability tests, to remove the confounding effects of gangue chemistry .

However, measuring the performance of a flotation tank containing a single-species pulp is difficult. For instance, the grade of the concentrate is undefined when there is no gangue in the system. However, by considering those behaviours and properties that single- and multi-species systems have in common, it is possible to obtain single-species proxies for metallurgical performance. After Norori-McCormac et al. (2017), the following proxies will be used for this work:

- concentrate flow rate as an indicator of recovery, and
- solids fraction as an indicator of grade.

For the system shown in equation 3, these can both be calculated from a sample of overflowing froth collected for a fixed amount of time at the lip of the tank. The concentrate flow rate is just the mass of concentrate collected at the lip, normalised to the total lip length, and the amount of time during which the sample was collected. The solids fraction is just the dry mass of the concentrate divided by the total wet mass collected.

3.3. Measuring froth stability

The stability of the froth has a significant effect on the performance of a flotation system. However, there is no single, widely-accepted definition of dynamic stability of an overflowing froth. A useful proxy for dynamic froth stability– proposed by Woodburn et al. (1994), and developed by Barbian et al. (2006) and others– is the the fraction of the air supplied to the flotation tank reporting to the concentrate as overflowing bubbles. That is, it is the fraction of the supplied air not lost from bubbles bursting on the surface of the froth.

This air recovery, α, can be estimated from the thickness of the overflowing layer of froth above the lip of the tank, h, and the speed at which it is overflowing, v, as follows Morrison (2017):

$$\alpha = \frac{\eta h v l}{Q} \tag{1}$$

where l is the overflowing lip length, η is the fraction of the froth that is air, and Q is the volume of air supplied to the tank per unit time.

Most recently, air recovery was shown to be a good proxy for froth stability in the laboratory-scale tank described herein, by Morrison 2017. All of the variables in equation 1, except for the air fraction of the froth– usually assumed to be unity– are measured using the instruments shown in figure 3. Therefore, since air recovery is so easily measured, and since it has been shown to correlate with froth flotation performance both at laboratory- and plant-scale (Hadler and Cilliers, 2009; Hadler et al., 2010), it is the measure of froth stability adopted here.

4. Experiment

4.1. Experimental design

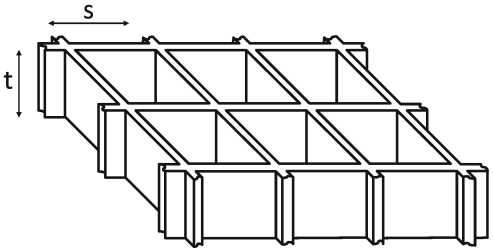


Figure 5: A schematic of the horizontal baffle design, showing pore size (s) and mesh thickness (t).

For ease of manufacture and parametrisation, a square lattice mesh design was chosen for this work. Figure 5 shows a portion of such a mesh, with its aperture size, **s**, and thickness, **t**, labelled. In addition to the pore size and mesh thickness, a third design parameter is included in this study: the mesh position, **p**, which is the distance between the top of the mesh and the pulp-froth interface.

Scoping tests were performed using the method of steepest ascent to determine the range of variability of each of these design variables. The resultant ranges of interest are given in table 1. Table 1 also shows a midpoint value, included in the experimental matrix in order to capture any quadratic behaviour in the performance response of the tank to changes in these three variables (Box et al., 2005). The result is a 3^3 experimental matrix.

Table 1: The three-level ranges of each of the mesh design variables.

Design Variable	**Levels (mm)**		
	-1	**0**	**+1**
Pore size	100	120	140
Thickness	20	40	60
Position	200	225	250

In order to produce a tractable experimental program, a Box-Behnken fractional factorial design was chosen (Box and Draper, 1959). This reduced the number of runs per experiment from the 30 that would be required for a full factorial experiment with repeated centre points, to just 15, without sacrificing the explanatory power of the resultant model. Importantly, a Box-Behnken experimental design does not compromise the explanatory power of cross-terms in such a model . This is particularly important for understanding the interplay of changes in design variables in a complex system such as this.

The Box-Behneken experimental design obtained in this way was randomised and run in triplicate on different days in order to isolate the effect of random environmental and system fluctuations on the resulting regression models. Each combination of conditions was run for 10 minutes to allow the system to settle, and then three overflowing froth samples were collected at minute intervals over the following 5 minutes. Corresponding average values of the froth depth and variability, the supplied air rate, and the overflowing froth height and speed– necessary to calculate the air recovery from the system– were calculated for a minute window centred at the time at which each froth sample was taken.

The result of this procedure was a matrix of corresponding indicative grade and recovery, and air recovery, for each combination of the predictor levels shown in table 1. This matrix formed the raw data from which a response surface model was obtained for each of measure of froth flotation tank performance discussed above.

4.2. Response surface modelling

Response surface modelling is a method of obtaining an *empirical* model of a response variable, such as the indicative grade, in terms of a range of predictor variables, such as the mesh design parameters listed in in table 1. The resulting model is descriptive rather than explanatory, and so cannot reliably be used to predict responses to combinations of predictor variable values outside the range of values for which the model is built. Furthermore, even the response surface bounded by the appropriate predictor ranges is only an approximation to the underlying true response, and is limited both by the model specification employed, and by the compounding uncertainties in the experimental measurement of predictors and responses. Nevertheless, a response surface model, based

on a well-designed experimental program, is a good balance between predictive power and model-building tractability (Myers et al., 2016).

Since a descriptive, rather than an explanatory, model is to be obtained here, there is no theoretical reason to include trigonometric or other functional forms in the model specification, and so any non-linearity will take the form of a polynomial. While it is, in some cases, necessary to perform a non-linear regression on a polynomial model specification, it is also possible– and simpler, and quicker– to perform a linear regression by linearising the polynomial. Linearising the polynomial is as simple as treating interaction, $x \cdot y$, and self-interaction, x^2, terms as distinct variables, rather than as being derived from the linear variables, x and y. Doing so significantly speeds up the model-fitting procedure, and allows a more robust model-fitting through stepwise regression.

Thus, for this work, a second-order model specification with pairwise interactions terms was chosen. It has the form:

$$\mathbf{R} = a_0 + a_1 \cdot p_1 + a_2 \cdot p_2 + ... + a_{12} \cdot p_1 \cdot p_2 + ... + a_{11} \cdot p_1^2 + a_{22} \cdot p_2^2 + ... \quad (2)$$

where $\mathbf{R}$ is the response, p_n are the predictors, and a_{nm} are the estimated model coefficients.

In the rest of this work, model specifications will be given in the concise notation of (Wilkinson and Rogers, 1973), distinguishable from algebraic notation by use of the tilde, $\sim$, instead of the equality, $=$, operator. In Wikinson-Rogers notation, the model coefficients are implied, and not given explicitly, so $\mathbf{R} = a_0 + a_1 \cdot p_1 + a_n \cdot p_1 \cdot p_n + a_{11} \cdot p_1^2$ in algebraic notation would be given as $\mathbf{R} \sim p_1 + p_1 : p_n + p_1 \wedge 2$ in Wilkinson-Rogers notation. Note also that cross terms are signified by a colon, and higher order terms by a caret followed by the order.

4.3. Predictor selection and model conditioning

Table 2: Symbols and descriptions of all of the predictors used to create response surface models for the mesh design experimental program, and their variance inflation factors (VIF) in terms of all other terms in the regression.

Symbol	**Name**	**Units**	**VIF**
s	Pore size	cm	2.358
t	Mesh thickness		3.282
p	Mesh position		5.528
Q	Supplied air rate	$\mathrm{cm} \cdot \mathrm{s}^{-1}$	3.249
Df	Froth depth	cm	2.071
dDf	Froth depth variability		5.785

Table 2 contains a summary of all of the predictors used to obtain response surface models of the performance of the laboratory-scale froth flotation tank. Regressions of the indicative grade and recovery, and air recovery, of the system performed with these predictors fulfil the requirements of model linearity, homoscedasticity, and multivariate normality required of a linear regression .

In addition to the mesh design parameters– **s**, **t**, and **p**– two tank operating variables are included in the model as predictors. They are the air flow rate into the tank, **Q**. and the depth of the pulp-froth interface below the overflowing lip of the tank, **Df**.

The supplied air rate is controlled by a proportional-integral-derivative (PID) control loop implemented in the computer control system. This control loop keeps the air flow rate within 2% of its setpoint value (Morrison, 2017), which– for the supplied air rate chosen for this work– means an air rate varying between 49 lpm and 51 lpm. While this is effectively a stable air supply, it was found that including the air rate as a predictor improved the goodness of fit of the resulting response surfaces, as measured by the coefficient of multiple determination, R^2, by between 5% and 10%.

The second operating variable included as a predictor, the froth depth, was included in two guises. Firstly, it was included in its own right to account for the change in froth depth due to the increased gas hold up brought about by the installation of horizontal meshes in the tank, and due to the randomly varying volume of supplied air rate. Doing so improved the coefficient of multiple determination of the regressions by between 10% and 20%, confirming that froth depth is a significant determinant of froth flotation performance .

Secondly, the variation of the froth depth over time, **dDf**, was also included as a predictor in the regressions. This variation was quantified as the standard deviation of the froth depth, **Df**, over a moving 10-second window. It is thus a proxy for the turbulence of the pulp-froth interface, and has been shown to be affected by variations in the supplied air rate, and in the impeller rotational speed (Morrison, 2017).

Also in table 2 are given the variance inflation factors (VIFs) of each of the predictors included in the regressions as a function of all of the others. These VIFs are a measure of the multicolinearity of the model, the degree of independence of the predictors. If a predictor were colinear, it would be impossible to isolate its effect on the response surface from the effects of the predictors upon which it depends. Kutner et al. (2004) proposed that a VIF greater than 4 requires some justification to include in a regression model, and that one with a VIF greater than 10 should be excluded. From table 2, it is clear that only the mesh position, **p**, and interface variability, **dDf**, show a degree of colinearity. However, neither show an unacceptable degree of colinearity, and so both are included as predictors; the mesh position because it is a design parameter under investigation in this work, and the interface variability because it is a proxy for the interface turbulence that the meshes are designed to reduce.

Further improvements in the goodness of fit of the resulting response surfaces were achieved by removing outlying experimental points that have leverage and influence beyond some threshold. The leverage of an outlier is a measure of how far its value lies from those of surrounding points. High-leverage outliers were removed following the procedure of Goodall (1993). The influence of an outlier can be quantified using its Cook's Distance (Cook and Weisberg, 1982). Influential outliers were also removed following the standard procedure.

4.4. Normalising the response

Response surface models obtained for the performance of the laboratory-scale flotation tank using the predictors in table 2 still suffered from one last problem: auto- or serial correlation. That is, the response to each randomised combination of predictors depended on the response to the preceding combination. In particular, this autocorrelation took the form of an exponential decay, suggesting that it arose from the decay of the surface hydrophobicity of the silica particles in the pulp over the course of an experiment. Including the elapsed time since the beginning of the experiment removed this autocorrelation. However, this was not the approach adopted for this work.

Instead, response surface models of the baseline performance of the unmodified laboratory-scale flotation tank were obtained as a function of tank operating conditions and experiment run time [Morrison 2017]. These baseline models were then used to obtain the predicted performance of the system under the conditions– other than the mesh design parameters– under which each run was performed. Subtracting the predicted meshless performance from the actual performance resulted in measures not of absolute performance, but the amount by which the meshed run over- or underperformed the prediction. Since the rebaselined performance of the modified tank thus included the effect of performance deterioration, as obtained in the baseline case, this removed the autocorrelation in the resultant models.

Finally, in order to simplify the interpretation of the regression statistics presented in this work, the value of both the responses and the predictors given in table 2 were normalised to the range $[0, 1]$. As a result, the performance of the modified tank is given as a fractional change relative to the baseline case, and the relative size of the estimated coefficient of a given predictor in the model is a measure of the relative importance of that predictor's effect on the model.

The result of the response surface modelling described above is a set of multivariate regression models for each of the laboratory-scale froth flotation tank performance indicators as a function of the mesh design parameters, the supplied air rate to the tank, the depth of the pulp-froth interface below the overflowing lip, and the interface variability– a proxy for interface turbulence. These can be used not only to elucidate the drivers of flotation performance under these conditions, but also to predict the performance response to different combinations of conditions within the range given in table 1.

5. Results

5.1. Metallurgical performance

Table 3 shows the result of performing a step-wise, multiple linear regression on the indicative grade and recovery from the froth flotation tank, as a function of the mesh design parameters and tank operating conditions given in table 2. The significance of the models are quoted as decimal logarithms of the usual p-values. In this form, the standard threshold for acceptable model significance is given as $log_{10}0.1 = -1$, and for good significance as $log_{10}0.05 \approx -1.3$. By this measure, both the indicative grade and recovery models are highly significant.

The goodness of fit of a regression can broadly be discussed in terms of how well the model accounts for variations in the training data upon which it is based, how significant the coefficient estimates are, and how well the model predicts new test data . The fraction of the variability in the training data that is accounted for by a response surface model is given by the adjusted coefficient of multiple determination, $\bar{R}^2$. The adjusted R^2 takes into account the number of predictors and sample points used in the final model, and thus avoids the large sample-large model bias that the ordinary R^2 statistic suffers from . The adjusted R^2 of the grade model is over 90%, and that of the recovery model is nearly 80% – they both account for a great majority of the variation in the underlying training data.

The threshold significance adopted for the step-wise linear regression was 10%, and so no higher order terms remain in the model whose significance is less than this. In table 3, significances are quoted as logarithms of the normal p-value, and so no higher order terms whose significance is less than $log_{10}p = -1$ appear in the table. However, less significant first order terms that appear in significant higher-order terms are retained in the model for completeness.

There are two formal measures of the significance of terms in the regression– a Students T-test as part of the regression procedure, and a Fishers F-test performed as part of an analysis of variance (ANOVA) of the resultant model [Hogg, 1992]. The F-test of the significance of an estimated coefficient has the null hypothesis that the addition of that predictor to the model has no effect on the predictive power of the model– technically, that it does not reduce the sum of the squares of the models residuals. Coefficients that are insignificant by this measure have little effect on the predictive power of the model, and can thus be safely omitted.

The T-test of the significance of an estimated coefficient has the null hypothesis that the coefficient is zero, that the predictor has no effect on the modelled response. The estimated coefficient of a given predictor in the regression model is a measure of the effect that changing that predictor by one unit has on the response. For a regression in which all of the predictor values have been normalised to the same range, a predictor coefficient that has a low significance, as measured by a T-test, will be relatively small compared to the other coefficients.

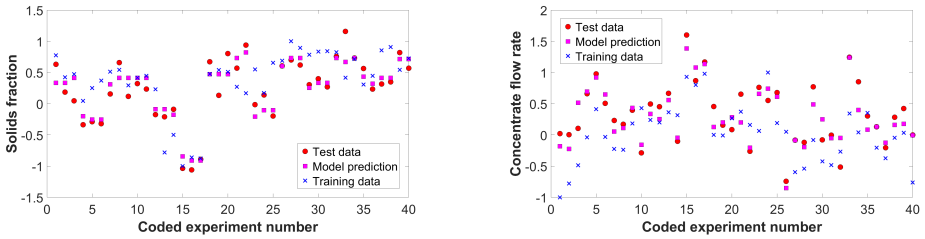
Table 3: Estimated model coefficients (Est. Coeff.) and their significance (Sig.) for each term in the grade (left) and recovery (right) regression models. The model and term significances are quoted as logarithms of the corresponding p-values, with threshold significances of $log_{10}0.05 \approx -1.3$ and $log_{10}0.1 = -1$.

	Indicative grade			**Indicative recovery**		
	Est. Coeff.	**Sig. ($\log_{10}$p)**		**Est. Coeff.**	**Sig. ($\log_{10}$p)**	
		Model	**ANOVA**		**Model**	**ANOVA**
Intercept	0.628	-6.921		0.836	-3.265	
s	0.012	-0.119	-2.688	0.222	-1.615	-0.363
t	-0.222	-4.756	-9.150	0.453	-4.886	-6.302
p	-0.031	-0.193	-3.421	0.639	-1.582	-0.113
Df				-0.174	-0.446	-3.670
dDf	-0.184	-1.450	-1.450	-0.068	-0.134	-1.325
Q	-0.089	-0.543	-0.831	0.083	-0.161	-0.274
s:t	0.470	-9.417	-9.417			
s:Df				-0.710	-3.614	-3.614
s:Q	0.581	-5.359	-5.359	-1.053	-2.923	-2.923
t:p	0.145	-1.542	-1.542	-0.389	-2.046	-2.046
t:Df				-0.423	-1.331	-1.331
t:Q	-0.377	-2.505	-2.505	0.885	-2.455	-2.455
p:Df				1.132	-3.922	-3.922
p:dDf				0.862	-1.980	-1.980
p:Q	0.486	-3.098	-3.098	-1.195	-1.897	-1.897
Df:dDf				0.920	-2.650	-2.650
s^2	-0.357	-6.125	-6.125			
t^2	-0.191	-2.453	-2.453	-0.473	-2.809	-2.809
p^2				-0.408	-1.171	-1.171
Q^2				-1.348	-4.577	-4.577
Goodness of fit ($\bar{R}^2$)			**0.916**			**0.795**
Model significance ($\log_{10}$p)			**-12.328**			**-4.324**

Thus, the coefficient estimate is a measure of the importance of a given predictor to the overall model. However, just because a given predictor has little effect on the modelled response to the training data, does not mean that it will not significantly affect future predictions.

Table 3 displays both the model and analysis of variance (ANOVA) significances of each term in the grade and recovery regression models. Those whose model or ANOVA p-values are insignificant are highlighted in orange. In all cases, the insignificant coefficients are first order terms, whose estimated coefficients are small compared to the intercept terms. In most cases, insignificant model p-values are complemented by significant ANOVA p-values, and vice versa, further supporting the retention of these terms in the model.

The exception to the above is the supplied air rate term, Q, which is insignificant both by the F- and the T-tests. However, two out of three of the most influential terms in both the grade and recovery models (highlighted in shades of blue) are higher order terms involving the supplied air rate. Thus, it is essential to retain this term in the model.



(a) Indicative grade

(b) Indicative recovery

Figure 6: A comparison of the model training (×) and predicted (■) data, with new test (●) data not included in the indicative grade (left) and recovery (right) regressions.

The final test of the goodness of fit of the grade and recovery regression models is how well they predict new test data. Figure 6 shows a comparison of the models with both the training data upon which they were based, and independent test data. Since the aim of this comparison is not to interpret the data, but just to see how well it is predicted and explained by the model, it is given in normalised units, and plotted in the order in which the experiments were conducted. Thus, any trend or structure in the data is coincidental.

From figure 6a, it is clear that the grade model both explains its training data, and predicts independent data very well. However, figure 6b does not support the same conclusions about the recovery model, which seems to consistently overestimate its training data. This is a physical representation of the lower adjusted coefficient of multiple determination of the indicative recovery compared to the indicative grade models, and of the larger number of insignificant recovery model terms. However, the indicative recovery model better predicts the independent test data, suggesting that the model is actually a fairly good

fit to the recovery behaviour of the system.

Thus, from the foregone discussion, it is clear that the indicative grade and recovery models obtained by response surface modelling of the metallurgical response of the froth flotation tank to varying mesh design parameters are both significant and a good fit to the underlying behaviour of the system.

Equations 3 and 4 are the indicative grade and recovery regression models in Wilkinson notation. The coefficients of each term are given in table 3. The three most significant terms in each model, corresponding to the blue highlights in table 3, are given in boldface.

$$\begin{aligned}\textbf{Grade} \sim 1 &+ \mathrm{dDf} \\ &+ (\mathbf{s} \cdot \mathbf{t}) + (\mathbf{s} \cdot \mathbf{Q}) + (\mathrm{t} \cdot \mathrm{p}) + (\mathrm{t} \cdot \mathrm{Q}) + (\mathbf{p} \cdot \mathbf{Q}) \\ &+ \mathrm{s\,\hat{}\,2} + \mathrm{t\,\hat{}\,2}\end{aligned} \tag{3}$$

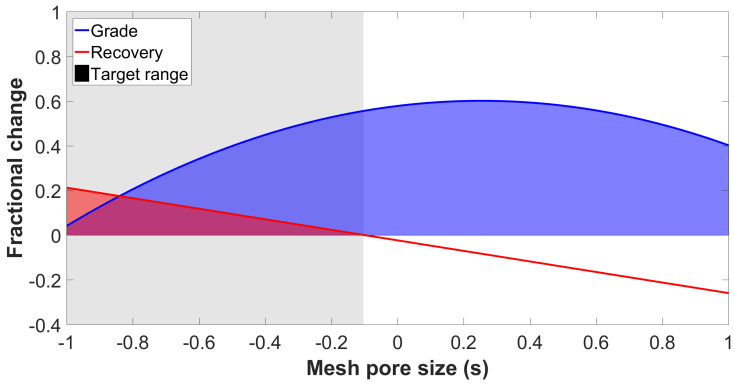
$$\begin{aligned}\textbf{Recovery} \sim 1 &+ (\mathrm{s} \cdot \mathrm{Df}) + (\mathrm{s} \cdot \mathrm{Q}) \\ &+ (\mathrm{t} \cdot \mathrm{p}) + (\mathrm{t} \cdot \mathrm{Df}) + (\mathrm{t} \cdot \mathrm{Q}) \\ &+ (\mathbf{p} \cdot \mathbf{Df}) + (\mathrm{p} \cdot \mathrm{dDf}) + (\mathbf{p} \cdot \mathbf{Q}) + (\mathrm{Df} \cdot \mathrm{dDf}) \\ &+ \mathrm{t\,\hat{}\,2} + \mathrm{p\,\hat{}\,2} + \mathbf{Q}\,\hat{}\,2\end{aligned} \tag{4}$$

From these equations, it is clear that indicative grade has a simpler response to the effects of varying mesh design parameters and tank operating conditions than does the indicative recovery, and that the depth and variability of the pulp-froth interface below the overflowing lip of the tank has both a larger, and a more complicated effect on the indicative recovery than grade.

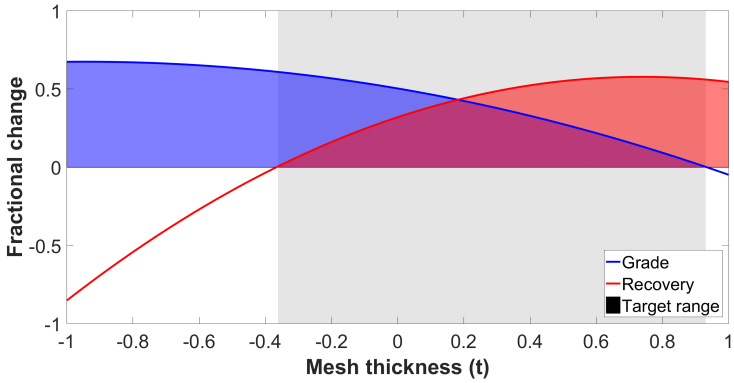
The marginal effect that a particular predictor, such as the mesh thickness, t, has on the indicative grade or recovery can be obtained by partially differentiating equation 3 or 4 by that predictor. Except in the simplest case of non-interacting predictors, such a marginal effect will still be in terms of other predictors. Thus, it is more convenient to visualise an aggregate effect, which is obtained by integrating the partially-differentiated marginal effect with respect to all of the predictors of which it is a function, over the full range of variability of those predictors. Thus, the overall aggregate effect of predictor x_1 on response R is given as

$$\overline{\partial_{x_1}\mathbf{R}} = \int\int ... \int dx_2 dx_3 ... dx_N \partial_{\delta x_1}\mathbf{R}(x_2, x_3, ..., x_N), \tag{5}$$

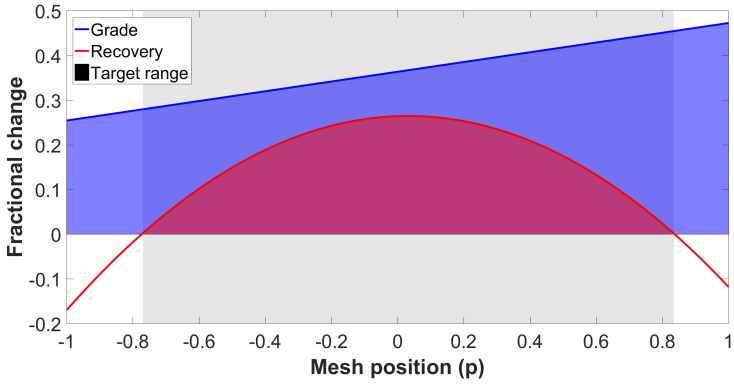
Figures 7 shows the aggregate effects of varying mesh design parameters on the indicative grade and recovery from the bench-scale froth floation tank. These indicative grades and recoveries have been rebaselined to the performance of the unmodified, meshless tank, and so are plotted as fractional changes relative to the unmodified case. Thus, positive fractional changes in figure 7 represent improvements in the grade or recovery relative to the unmodified case.



(a) Normalised mesh pore size.



(b) Normalised mesh thickness.



(c) Normalised mesh position.

Figure 7: A summary of the fractional changes in indicative grade and recovery as a function of each of the mesh design variables.

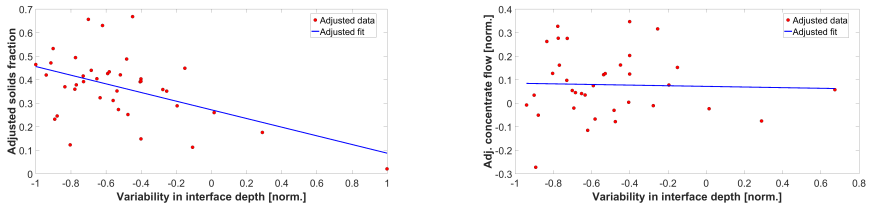


Figure 8: The aggregate effect on indicative grade (left) and recovery (right) of the variability of the pulp-froth interface.

In some cases, the indicative grade and recovery curves in figure 7 suggest optimal mesh design parameters that maximise either grade or recovery. Since froth flotation grade and recovery are complementary , these do not coincide. However, there exist target ranges (shaded in grey in figure 7) for each design parameter over which both the indicative grade and recovery exhibit improvements over the unmeshed case. These target ranges provide the constraints for a bounded, multivariate optimisation of the grade and recovery in terms of the mesh design parameters.

5.2. Froth stability

While the response of the metallurgical performance of the froth flotation tank to changes in mesh design parameters exhibits the classic trade-off behaviour between indicative grade and recovery, its response to changes in the variability of the pulp-froth interface does not. Instead, as demonstrated by the aggregate responses in figure 8, the lower the interface variability, the higher both the resultant indicative grade and recovery. This is most pronounced in the case of the grade, which is more sensitive to the froth stability, which declines with increasing interface variability [Morrison 2017]. Thus, it is instructive to consider the effect of mesh design on the pulp-froth interface.

Table 4 shows the result of performing a step-wise, multiple linear regression on this interface variability as a function of mesh design parameters. As in table 3, significances are quoted as decimal logarithms, and the same significance thresholds are adopted. While both the overall pulp-froth interface variability model is highly significant, its coefficient of multiple determination is about 71%– lower than either the indicative grade or recovery models given in table 3. This lower goodness of fit is also reflected in the larger number of insignificant terms in these models, compared to the grade and recovery models. Nevertheless, as shown in figure 9, the interface variability model does at least as good a job explaining its training data and predicting new data as does the indicative recovery model, and so can be afforded as much credibility.

Equation 6 shows the pulp-froth interface variability model in Wilkinson notation, with the three most influential terms (highlighted in blue in table 4) in boldface. Like the indicative recovery model, the interface variability model is

Table 4: Estimated model coefficients (Est. Coeff.) and their significance (Sig.) for each term in the pulp-froth interface variability regression model. The model and term significances are quoted as in table 3.

	Interface variability		
	Est.	**Sig. ($\log_{10}$p)**	
	Coeff.	**Model**	**ANOVA**
Intercept	-1.258	-10.235	
s	-0.281	-2.621	-0.378
t	-0.071	-0.695	-2.741
p	-0.132	-0.955	-2.992
Df	0.282	-1.782	-2.417
Q	-0.179	-0.408	-2.005
s:t	-0.146	-1.106	-1.106
s:Df	0.539	-3.769	-3.769
s:Q	0.299	-1.087	-1.087
t:Df	0.373	-1.617	-1.617
t:Q	-1.188	-4.334	-4.334
t^2	0.404	-4.679	-4.679
p^2	0.451	-2.021	-2.021
Df^2	0.364	-2.054	-2.054
Q^2	0.761	-2.432	-2.432
Goodness of fit ($\bar{R}^2$)			**0.712**
Model significance ($\log_{10}$p)			**-4.450**

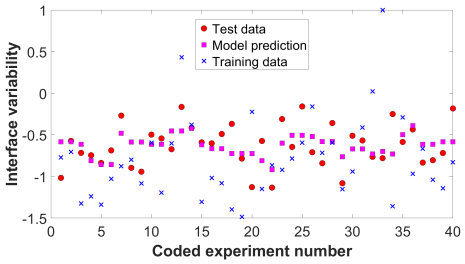


Figure 9: A comparison of the model training (×) and predicted (■) data, with new test (●) data not included in the pulp-froth interface variability regression.

fairly complicated, and dominated by cross-terms involving the flow rate of air supplied to the tank, Q.

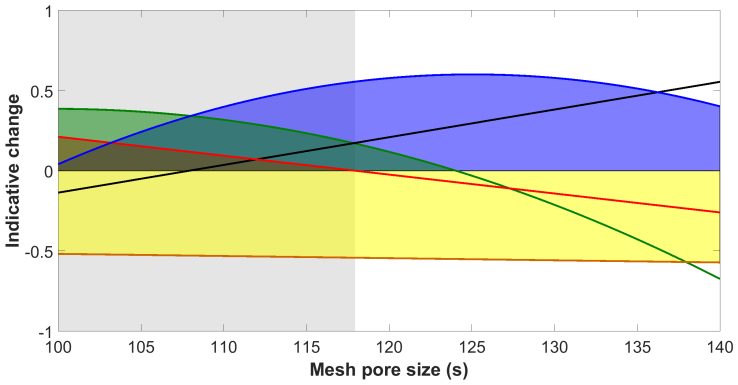
$$\begin{aligned}\textbf{Variability} \sim 1 + p + (\mathrm{s}\cdot\mathrm{t}) + (\mathbf{s}\cdot\mathbf{Df}) + (\mathrm{s}\cdot\mathrm{Q}) \\ + (\mathrm{t}\cdot\mathrm{Df}) + (\mathbf{t}\cdot\mathbf{Q}) \\ + \mathrm{t}\,\hat{}\,2 + \mathrm{p}\,\hat{}\,2 + \mathrm{Df}\,\hat{}\,2 + \mathbf{Q}\,\hat{}\,2\end{aligned} \tag{6}$$

Figure 10 shows the aggregate responses of the indicative grade and recovery, and pulp-froth interface variability to variations in the mesh design variables. As in figure 7, the range of values for which the changes in both the indicative grade and recovery relative to the meshless case are positive are highlighted in grey. These figures show that, except in the case of the the mesh pore size, for which the indicative recovery has no maximum, the pulp-froth interface variability is minimised where the combined change in indicative grade and recovery is maximised. Furthermore, this correlation is most pronounced between the indicative recovery and pulp-froth interface variability.

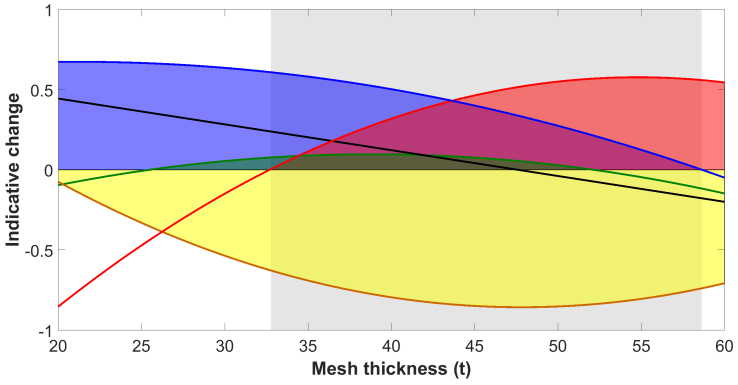
6. Conclusions

The foregone has shown that installing a horizontal mesh in a laboratory-scale froth flotation tank, by reducing the turbulence at the pulp-froth interface, allows both the indicative grade and recovery achievable to be improved relative to the unmodified tank run under otherwise similar conditions.

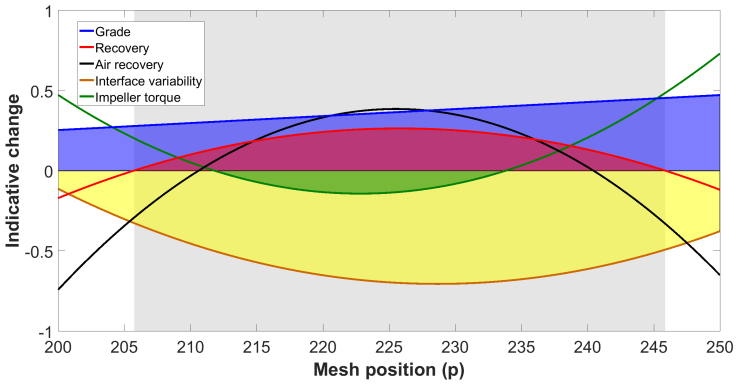
The ranges of mesh design parameters over which both the indicative grade and recovery show improvement, represent the bounds of a constrained, multivariate optimisation problem that results in a grade-recovery curve that can be used to design retrofittable meshes to maximise either grade or recovery.



(a) Normalised mesh pore size.



(b) Normalised mesh thickness.



(c) Normalised mesh position.

Figure 10: A summary of the fractional changes in indicative grade and recovery, and pulp-froth interface variability, as a function of each of the mesh design variables.

7. References

N Barbian, K Hadler, and JJ Cilliers. The froth stability column: Measuring froth stability at an industrial scale. *Minerals Engineering*, 19(6):713–718, 2006. ISSN 08926875. doi: 10.1016/j.mineng.2005.09.021.

George E. P. Box, J. Stuart Hunter, and William Gordon Hunter. *Statistics for experimenters : design, innovation, and discovery*. Wiley-Interscience, 2005. ISBN 0471718130.

GEP Box and NR Draper. A Basis for the Selection of a Response Surface Design. *Journal of the American Statistical Association*, 54(287):622–654, sep 1959. ISSN 0162-1459. doi: 10.1080/01621459.1959.10501525.

RD Cook and S Weisberg. *Residuals and influence in regression*. Chapman and Hall, 1982. ISBN 041224280X.

T Eisele, SJ Ripke, and SK Kawatra. Stabilization of Flotation Column Performance by Horizontal Baffle Columns. In BK Parekh and JD Miller, editors, *Advances in flotation technology*, pages 391–396. Society for Mining, Metallurgy, and Exploration, Littleton, CO, 1999. ISBN 0873351843.

CR Goodall. Computation using the QR decomposition. *Handbook of Statistics*, 9:467–508, 1993. doi: 10.1016/S0169-7161(05)80137-3.

F. Gourram-Badri, P. Conil, and G. Morizot. Measurements of selectivity due to coalescence between two mineralized bubbles and characterization of MIBC action on froth flotation. *International Journal of Mineral Processing*, 51(1): 197–208, 1997. ISSN 03017516. doi: 10.1016/S0301-7516(97)00034-3.

A Gupta and DS Yan. *Mineral processing design and operation : an introduction*. Elsevier, 2006. ISBN 9780444516367.

K Hadler and JJ Cilliers. The relationship between the peak in air recovery and flotation bank performance. *Minerals Engineering*, 22(5):451–455, 2009. ISSN 08926875. doi: 10.1016/j.mineng.2008.12.004.

K Hadler, CD Smith, and JJ Cilliers. Recovery vs. mass pull: The link to air recovery. *Minerals Engineering*, 23(11):994–1002, 2010. ISSN 08926875. doi: 10.1016/j.mineng.2010.04.007.

GJ Jameson. The effect of surface liberation and particle size on flotation rate constants. *Minerals Engineering*, 36:132–137, 2012. ISSN 08926875. doi: 10.1016/j.mineng.2012.03.011.

S Kawatra and T Eisele. Flotation column with adjustable supported baffles, 1994.

MH Kutner, C Nachtsheim, and J Neter. *Applied linear regression models*. McGraw-Hill/Irwin, Boston, New York, 4th ed. edition, 2004. ISBN 9780072386912.

Angus James Morrison. *Tank design modifications for the improved performance of froth flotation equipment*. PhD thesis, Imperial College of Science, Technology and Medecine, 2017.

RH Myers, DC Montgomery, and CM Anderson-Cook. *Response surface methodology: process and product optimization using designed experiments.* Wiley, 2016. ISBN 9781118916032.

T Napier-Munn. *Statistical methods for mineral engineers : how to design experiments and analyse data*. Julius Kruttschnitt Mineral Research Centre., 2014. ISBN 0980362245.

A Norori-McCormac, PR Brito-Parada, K Hadler, K Cole, and JJ Cilliers. The interaction between particle size and air rate and the effect on froth stability. *Separation and Purification Technology*, 184:240–247, 2017.

MJ Prince and HW Blanch. Bubble coalescence and break-up in air-sparged bubble columns. *AIChE Journal*, 36(10):1485–1499, oct 1990. ISSN 0001-1541. doi: 10.1002/aic.690361004.

John Ralston, Daniel Fornasiero, and Robert Hayes. Bubbleparticle attachment and detachment in flotation. *International Journal of Mineral Processing*, 56 (1):133–164, 1999. ISSN 03017516. doi: 10.1016/S0301-7516(98)00046-5.

GN Wilkinson and CE Rogers. Symbolic Description of Factorial Models for Analysis of Variance. *Applied Statistics*, 22(3):392, 1973. ISSN 00359254. doi: 10.2307/2346786.

BA Wills. *Mineral processing technology : an introduction to the practical aspects of ore treatment and mineral recovery*. Butterworth-Heinemann, 2006. ISBN 9780750644501.

ET Woodburn, LG Austin, and JB Stockton. Froth based flotation kinetic model. *Chemical Engineering Research and Design*, 72(A2), 1994.