Mine stockpile design to minimise environmental impact

J. E. Everett
Faculty of Economics and Commerce,
The University of Western Australia

Abstract

In Western Australia, iron ore is extracted from inland mines and railed to a port, where it is loaded onto ships for export. Quality depends upon uniform composition ("grade control"), not only in iron, but also in several contaminant minerals. To achieve grade control, and to provide a buffer between production and demand, crushed ore is stored on to large stockpiles and then reclaimed, either at the mine or at the port. Environmental impact (and cost) is reduced if the land area devoted to stockpiles can be reduced without loss of grade control. There may also be environmental benefits in building the stockpiles at the mine rather than at the port. The stockpile array can be considered as a low-pass filter, filtering out short-term fluctuations in composition. Techniques considered include the use of multiple build and/or reclaim stockpiles, with and without intelligent stacking and reclaiming. Because of the complex correlations between the minerals and across time, design of the stockpile array is not readily amenable to mathematical analysis. Simulation using historic data provides a useful insight for comparing available stockpile policies. A simulation model has been constructed to aid the mine planner choose between alternative stockpile policies. The model has been constructed using VBA in Excel, to provide ease of use and portability.
Keywords: quality control, decision support, mining, stockpile design, intelligent stacking/reclaiming, forecasting, simulation.

1 Introduction

Iron ore from inland mines is railed to port, where it is crushed, separated into lump and fines products and stored onto large stockpiles. Ships are loaded for export with ore reclaimed from the stockpiles.

Development and Application of Computer Techniques to Environmental Studies X
G. Latini, G. Passerini & C. A. Brebbia (Eds) , www.witpress.com, ISBN 1-85312-718-3

The exported ore is destined for blast furnaces, whose efficient operation depends upon the ore being of consistent composition, not only in iron but also in the major contaminants, including silica, alumina and phosphorus. The quality of the ore therefore depends upon each shipment closely matching the target composition vector.

Putting the ore onto stockpiles before shipment serves a double purpose. The stockpiles provide a quantity buffer between the arrival of ore trains from the mine and the more irregular departure of ships. But of just as much importance is the effect of stockpiles in blending out variations in ore composition.

Typically, stockpiles are of the order of 100 kilotonnes when full, with a footprint about 30 metres wide and 100 metres long. Stockpiles are built in a "chevron" pattern, with the stacker running back and forth the whole length of the pile. The stockpiles are built to completion, so that reclaiming does not start until stacking is finished. Reclaiming is from one end, in three "pilgrim steps", so any shipment contains closely equal portions taken from the top, middle and bottom of the pile.

The method of stacking and reclaiming is designed to ensure that the ore stacked to and then reclaimed from a stockpile is effectively uniform in composition. Studies comparing assays of stacked ore and reclaimed ore have shown that this goal is achieved, at least to within the errors of assay measurement. Further improvement in quality control therefore depends on controlling the variability in composition between successive stockpiles.

The port yard comprises an array of stockpile sites. If ore is stacked on arrival to a single stockpile, which is built to completion, then the larger the stockpile the more composition variability is smoothed out.

There may also be a quality benefit in having multiple completed stockpiles available at any one time, so further quality improvement can be achieved by loading ships with ore blended from multiple stockpiles, but the present paper does not consider this.

Increasing the size and number of stockpile sites comes at considerable economic and environmental cost. Coastal land in the vicinity of the port is mainly mangrove swamp, part of a complex and sensitive ecosystem. It is desirable to minimise the disturbance to the ecosystem by designing a stockpile system able to give adequate quality control with as small a stockpile footprint as possible.

Pitard [1] provides a thorough study of stockpile procedures for mining. Everett [2,3] discusses methods of iron ore quality control. Kamperman et al [4] describe application of these methods to a West Australian iron ore mining operation.

2 Scope

This paper describes a computer-based simulation model to aid the design of a stockpile system, based upon real data, to give improved quality control while minimising the stockpile footprint. The model design and development fully

Development and Application of Computer Techniques to Environmental Studies X
G. Latini, G. Passerini & C. A. Brebbia (Eds) © 2004 WIT Press, www.witpress.com, ISBN 1-85312-718-3

involved the mine operators and planners, both to incorporate their domain knowledge and to gain the cooperation and enthusiasm. To provide them with a "fishing line rather than a fish", a simulation model was built which they could use to explore alternative scenarios. To avoid software licensing fees and the costs of user acceptance and training, a dedicated simulation package was not used. Instead, the simulation model was built using Visual Basic (VBA) macros in Excel. Since staff were accustomed to using Excel spreadsheets, they quickly felt comfortable using a simulation model that had the appearance of an Excel workbook, with menus and buttons to drive the macros.

3 Quality control

Iron ore quality requires each shipment of ore to match target composition as closely as possible. The target composition is a vector of four components {iron, silica, alumina, phosphorus}, so there are four objective functions, one for each mineral. The customer's objective function may be skewed: for example, they may tolerate iron above target or silica and alumina below target more happily than the opposite. But, for the producer, to deliver ore richer than target is an opportunity cost. So we must treat the objective function as symmetric.

For each mineral, we can identify a tolerance level. If the composition departs from target by one tolerance unit in any mineral, the customer is equally unhappy. We can therefore define a dimensionless stress in each mineral, as the departure from target divided by the tolerance. Squaring the component stresses and adding then together defines a total stress, which provides a combined objective function to be minimised. This quadratic objective function, together with a set of linear constraints, is used to select the ore mined from the available sources within the open-cut mine pit. It is a classic quadratic programming problem (see Taha [5]), although for convenience we use the heuristic optimiser "Solver", supplied with Excel, for making the daily mining decisions.

Variations in the composition of the mined ore have long, medium and short-term frequency components, which we need to filter out. Monthly mine plans are designed to remove long-term variations. The daily mine plan should smooth medium-term variations. Building and reclaiming stockpiles, the subject of this study, serves to filter out short-term variations, with period up to the capacity of the stockpile yard.

The monthly mine plan, daily schedule and stockpile building decisions are all further complicated by the problem that we do not know the "true" composition of any ore until it has been mined, railed to port, crushed, sampled and assayed. All the operational decisions as to the mining and stockpiling of ore have to be made before its assay results are known. The mine planning decisions have to be made using coarsely spaced data from exploratory drill holes. Daily scheduling decisions can use more detailed estimates from assays of material collected during the drilling of blast holes prior to the ore blocks being broken up ready for mining. Even the blast hole estimates have considerable random error and bias compared to the final assays from the mined and crushed ore.

Development and Application of Computer Techniques to Environmental Studies X
G. Latini, G. Passerini & C. A. Brebbia (Eds) © 2004 WIT Press, www.witpress.com, ISBN 1-85312-718-3

To improve the information for daily scheduling and stockpiling decisions, we adjust the blast-hole estimates by a regression model, and add a bias estimated from a forecast based on the exponentially smoothed difference between the port assays and the blast hole estimates of recently mined ore. Hanke and Reitsch [6] discuss forecasting methods using exponential smoothing.

4 Building stockpiles

4.1 Assay Variability

We have seen that, within assay measurement error, ore stacked to and recovered from a stockpile can be considered uniform. Our goal is to minimise the variation in composition between stockpiles.

The assay data unit is a trainload (or more strictly a rake, or half-train) of about 10 kilotonnes. Trainloads are stacked to build stockpiles that can be from about 50 to 200 kilotonnes. If the composition of successive trainloads were uncorrelated, and stockpiles were built to completion from successive trainloads, then the variation in stockpile composition would be inversely proportional to the square root of the stockpile size.

However, not only do the assay data show strong correlations between minerals (silica and alumina are positively correlated with each other, and negatively with iron), but they also have strong serial correlations, usually positive. Because of the serial correlation, variability between stockpiles falls off more slowly than the inverse square root of stockpile size.

The complicated correlations, between minerals and across time, make it unwise to do any planning based upon synthetic data. Any modelling we do is based on historical data. Planning for a new operation uses data from some other operation that is geologically matched as closely as possible.

4.2 Stacking policies

Our purpose is to design a stockpile system that gives adequate control of variability while using as small a total stockpile footprint as possible.

Building larger stockpiles can improve the variability, but at the cost of increasing the stockpile footprint. For example, to halve the assay variability would require that the stockpile footprint be increased more than four-fold.

As an alternative to increasing stockpile size, we can stack each trainload to a destination chosen from multiple current stockpiles, as shown in Figure 1. The choice of stockpile can be random, sequential or "intelligent".

The incoming trainload has weight w and composition stress vector $\underline{x}$. The stress vector has components $\{x_k\}$, for minerals k=1 .. K. Each component of the stress vector is defined as the departure from target divided by the tolerance for that mineral.

The n^{th} stockpile (n=1,2, ... N) has weight W_n, and the stress vector $\underline{X}_n=\{X_{nk}\}$.

Development and Application of Computer Techniques to Environmental Studies X
G. Latini, G. Passerini & C. A. Brebbia (Eds) © 2004 WIT Press, www.witpress.com, ISBN 1-85312-718-3

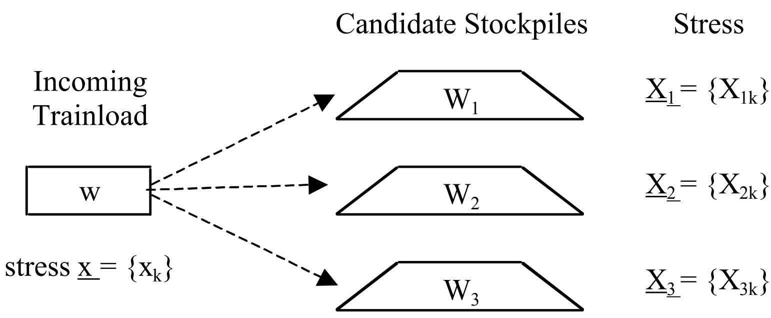

Figure 1: Stacking to multiple candidate stockpiles.

4.3 Sequential or random stacking

With sequential stacking, each trainload is sent in turn to each of the stockpiles currently being built. Random stacking is similar, but the destination choice is either truly random, or based on operational convenience (such as interaction of the stacker with other equipment).

If the composition of successive trainloads has no serial correlation, sequential or random stacking provides no benefit over stacking a single stockpile to completion. However, if they are serially correlated (equivalent to periodic variation), then sequential or random stacking removes some of the periodic variation by spreading each stockpile over a longer production period.

4.4 Intelligent stacking

Intelligent stacking uses the composition of the incoming ore, or an estimate of its composition, to decide the destination stockpile.

If we know the tonnages and estimated compositions of the incoming trainload and of the candidate stockpiles being built, we can send the trainload to the destination chosen to minimise the total stress of the stockpile yard.

As in Figure 1, the incoming trainload has weight w and composition stress vector $\underline{x}$. The stress vector has components $\{x_k\}$, for minerals k=1 .. K. The total weighted stress "s" of the trainload is given by:

$$s = w\underline{x}.w\underline{x} = w^2\underline{x}.\underline{x} = w^2\sum x_k^2 \tag{1}$$

Similarly, the nth stockpile, with weight W_n and stress vector $\underline{X}_n$, has weighted stress $S_n = W_n^2\underline{X}_n.\underline{X}_n$. So the current total weighted stress of the stockpiles is:

$$\sum S_n = \sum W_n^2\underline{X}_n.\underline{X}_n = \sum W_n^2\sum X_{nk}^2 \tag{2}$$

If we add the trainload to the nth stockpile, its weighted stress will increase by:

G. Latini, G. Passerini & C. A. Brebbia (Eds) © 2004 WIT Press, www.witpress.com, ISBN 1-85312-718-3

$$\delta S_n = (w+W_n)^2\sum[(wx_k+W_nX_{nk})/(w+W_n)]^2 - W_n{}^2\sum X_{nk}{}^2$$

$$= 2wW_n\sum x_kX_{nk}+w^2\sum x_k{}^2 = 2wW_n\underline{x}\mathbf{.}\underline{X}_n + w^2\sum x_k{}^2 \qquad (3)$$

Thus, to minimise the resulting total weighted stress, the trainload should be sent to the stockpile giving the smallest value for $W_n\underline{x}\mathbf{.}\underline{X}_n$ (the stockpile weight multiplied by the scalar product of the trainload and stockpile stress vectors).

4.5 Estimating train composition

Intelligent stacking, as described above, depends upon knowledge of the composition of the incoming ore. As discussed, at the time the stacking decision has to be made the assay results for the trainload are not yet available.

So in place of the true stress vector $\underline{x}$ we have to use an estimated $\underline{x}^*$. For intelligent stacking, we now send the trainload sent to the stockpile giving the smallest value for $W_n\underline{x}^*\mathbf{.}\underline{X}_n$

To estimate $\underline{x}^*$, we take the blast hole estimates of composition for the mined ore that went on to the train, regress these estimates towards the mean and then add a bias. The bias is the exponentially smoothed error of recent assays. The regression and exponential smoothing coefficients are established by analysis of previous blast hole and port assays, and reviewed every few months.

In the limit, if the estimated stress vector $\underline{x}^*$ is uncorrelated with the true stress $\underline{x}$, intelligent stacking becomes random stacking. In the simulation studies to be described, we investigated the effect of progressively blurring the estimated stress, from perfect knowledge, through the accuracy of available estimates, down to purely random assignment.

4.6 Stockpile size, number and footprint

Our goal is to limit the area devoted to stockpiles, while still filtering out composition variation for ore stacked and reclaimed through the stockpile array.

Stockpiles are built to a specific height and width, constrained by the stacking and reclaiming equipment. The length of a stockpile footprint is therefore proportionate to the stockpile capacity, and a partly built stockpile still needs to have the full footprint dedicated to it. Ignoring the comparatively small space lost between stockpiles, the area devoted to stockpiles is proportional to the number of stockpile sites, multiplied by the tonnage capacity of each stockpile.

Because stockpiles are built to completion before reclaiming, there will always be one stockpile currently being reclaimed. To provide a buffer between ship demand and trainload supply, a further stockpile site is needed, either full, waiting to be reclaimed, or empty, waiting to be stacked. So, if stacking is to N multiple sites at any time, then N+2 stockpile sites need to be made available, and the total stockpile footprint will be approximately proportional to (N+2)C, where "C" is the capacity of each stockpile site.

If intelligent reclaiming is being considered, the number of stockpile sites must be further increased to allow for multiple reclaim piles. However, intelligent reclaiming is beyond the scope of the present paper.

Development and Application of Computer Techniques to Environmental Studies X
G. Latini, G. Passerini & C. A. Brebbia (Eds) © 2004 WIT Press, www.witpress.com, ISBN 1-85312-718-3

5 The simulation model

A simulation model was built, to help the operators and planners to explore implications of the alternative stacking policies. The model was designed to help them develop stacking policies to give adequate control of variability while minimising the footprint of the stockpile array.

To make the model accessible and easily used by the staff, it was designed as an Excel workbook, with the simulation driven by macros. The specification sheet is shown in Figure 2. It allows the user to specify parameters for the desired simulation run.

Run Name	**10 kt Rakes (Mine 1, based on Mine 2 7/9/01-13/11/03)**							
Data Specification			Tolerance		Modify Input Rakes to		Correln	Rakes
			Stress	Ship	Mean	Std Devn	Port/Blast	Stress
Rake kt	10	Fe	1.00	0.40	62.15	0.96	0.67	0.96
Wt Index	1	P	.0200	.0050	.0549	.0054	0.64	0.27
Replications	1	SiO_2	0.33	0.31	3.41	0.64	0.57	1.93
		Al_2O_3	0.16	0.14	1.92	0.45	0.77	2.83

Single Run Simulation

Stack Piles	3	Pile kt	90		
Rakes/Pile	9	Footprint kt	450		
		3	Piles		Pile/
Stack Build Policy			S.D.	Stress	Rake
		Fe	0.16	0.16	17%
○ Sequential		P	.0014	0.07	27%
○ Random		SiO_2	0.15	0.45	23%
● Intelligent		Al_2O_3	0.08	0.52	18%
○ Perf. Knowl.		Total		0.71	20%

Multiple Run

Choose up to	Stack Piles					
Five Plots 1	1	2	3	4	5	6
To Completion	☒					
Sequential		☐	☐	☐	☐	☐
Random		☐	☐	☐	☐	☐
Intelligent		☒	☒	☐	☐	☐
Perf. Knowledge		☐	☐	☐	☐	☐

X-Axis, kt	X-Axis Range (kt)		
	Min	Max	Step
● Pile			
○ Footprint	90	210	30

Figure 2: Simulation model specification sheet.

The simulation uses data, stored in a sheet of the workbook. In the example here, the data are assays for 1800 trainloads, as actually dumped at the port. The statistical properties of the historic data can be adjusted to match a desired scenario, and blast hole estimates are generated, of accuracy determined by the specified correlation between port and blast.

G. Latini, G. Passerini & C. A. Brebbia (Eds) © 2004 WIT Press, www.witpress.com, ISBN 1-85312-718-3

Pressing either the “Single Run Simulation” or the “Multiple Run” button initiates the VBA macros that run the simulation.

5.1 Single Run

The Single Run option simulates a single scenario and reports the composition of the resulting completed stockpiles, in tabular and graphic form.

As an example, Figure 3 shows the silica composition of completed stockpiles, using intelligent stacking to three stockpile destinations. We see that six of the 200 completed stockpiles lie outside the tolerance range for silica.

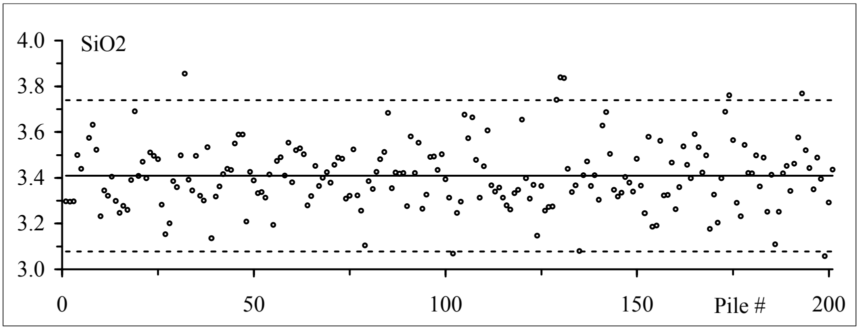

Figure 3: Silica composition for completed stockpiles.

Stacking and reclaiming ore through a stockpile is a low-pass filter, filtering out short-period variations in composition.

The simulation reports Fourier analyses to explore the filtering. The variance is cumulated from zero frequency, and expressed as a cumulative standard deviation, for input trainloads and for completed stockpiles.

The frequency analysis graphs for silica are shown in Figure 4. The graph for each cumulated standard deviation stops at the relevant Nyquist (or folding) frequency.

The Nyquist frequency has period twice the data interval, so has period 20 kt for 10 kt trainloads and 180 kt for 90 kt stockpiles.

If stacking to a single stockpile, the cumulated standard deviations would overlay, and the reduction (from 0.64 to 0.30) achieved by stockpiling would be equivalent to stopping the graph at the stockpile Nyquist frequency.

The further reduction in standard deviation (from 0.30 to 0.15) is achieved by intelligent stacking.

5.2 Multiple Run

A Multiple Run simulation explores and compares a set of scenarios. Each scenario is specified as to the number of stack piles, and whether they are to be built in sequence, at random, or with intelligent stacking. The unattainable scenario of perfect knowledge, with assay values known accurately, may also be simulated, to provide a comparison.

Development and Application of Computer Techniques to Environmental Studies X
G. Latini, G. Passerini & C. A. Brebbia (Eds) © 2004 WIT Press, www.witpress.com, ISBN 1-85312-718-3

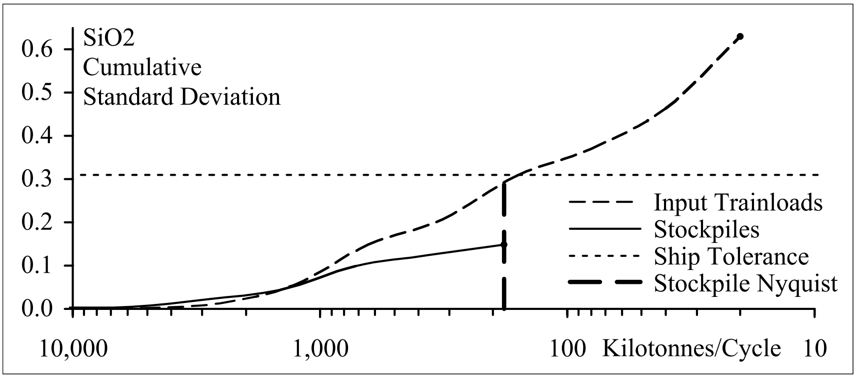

Figure 4: Frequency analysis.

For each scenario, the simulation is repeated for a series of stockpile sizes. The standard deviations are reported, tabulated and as graphs, plotted against either stockpile size or total footprint. Figure 5 shows the standard deviation for silica plotted against stockpile size, for stacking to a single stockpile destination, and for intelligent stacking to two and to three stockpile destinations.

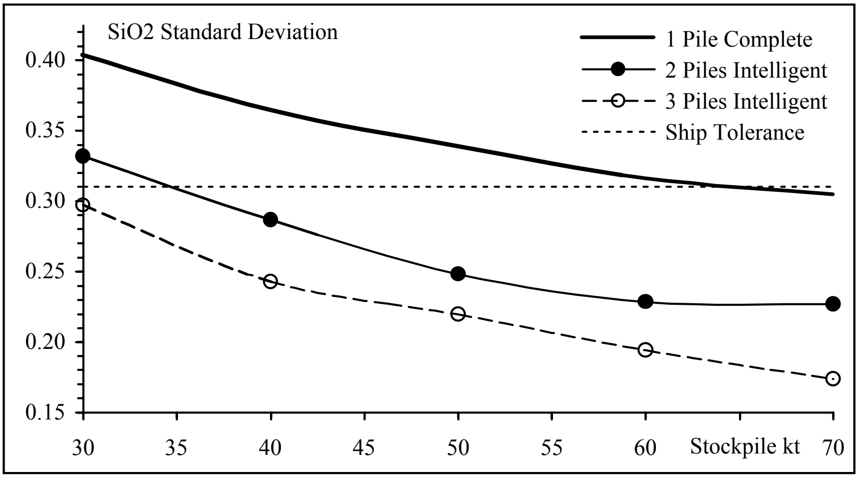

Figure 5: Intelligent stacking plotted against stockpile size.

Figure 5 suggests intelligent stacking to two and to three stockpiles yields increasing benefit. But increasing the number of stockpiles increases the footprint. Figure 6 plots the results against footprint. We now see intelligent stacking to two stockpiles gives improved control for a given footprint size, but no further benefit is achieved by intelligent stacking to three stockpiles.

G. Latini, G. Passerini & C. A. Brebbia (Eds) © 2004 WIT Press, www.witpress.com, ISBN 1-85312-718-3

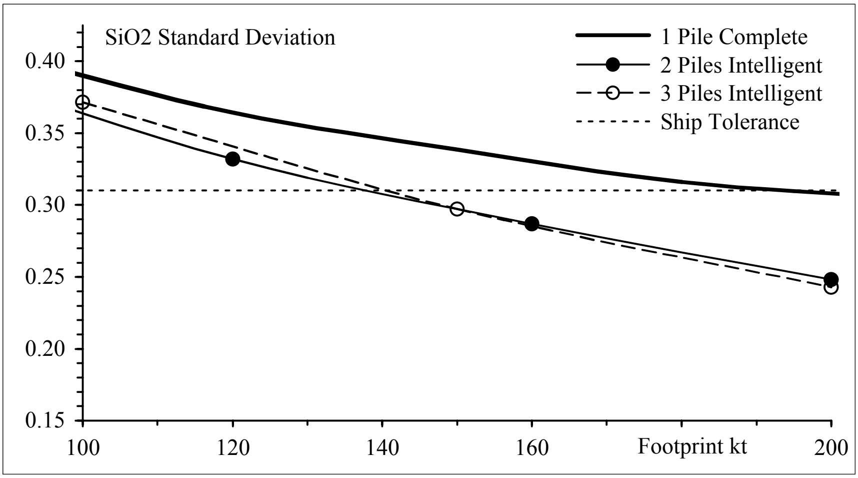

Figure 6: Intelligent stacking plotted against stockpile footprint.

6 Conclusions

We have seen that an Excel-based simulation model can help practitioners explore scenarios for stacking ore, so as to reduce the area used for stockpiles while maintaining control over variability in composition of the shipped ore. The modelling can be readily extended to explore alternative reclaiming scenarios.

References

[1] Pitard, F.P. *Pierre Gy's Sampling Theory and Sampling Practice,* 2nd ed., CRC Press: Boca Raton, Florida, 1993.

[2] Everett, J.E. Iron ore handling procedures enhance export quality, *Interfaces*, **26(6)**, pp. 82-94, 1996.

[3] Everett J.E. Iron ore production scheduling to improve product quality, *European Journal of Operational Research*, **129**, pp. 355-361, 2001.

[4] Kamperman, M., Howard, T. & Everett, J.E., Controlling product quality at high production rates. *Proc. of the Iron Ore 2002 Conference,* ed. R. Holmes, Australasian Institute of Mining and Metallurgy: Carlton, Victoria, pp. 255-260, 2002.

[5] Taha, H.A., *Operations Research,* 4th ed., Collier MacMillan: London, 1987.

[6] Hanke, J.E. and Reitsch, A.G., *Business Forecasting,* 6th ed., Prentice-Hall: Englewood Cliffs, 1998.

G. Latini, G. Passerini & C. A. Brebbia (Eds) © 2004 WIT Press, www.witpress.com, ISBN 1-85312-718-3