

Simulation Modeling of an Iron Ore Operation to Enable Informed Planning

J. E. Everett

University of Western Australia, Nedlands, WA, Australia

Jim.Everett@uwa.edu.au

Abstract

An iron ore mining company, operating in Western Australia, trucks ore from three geographically isolated sources to a crusher, where it is blended before and during crushing. The company prides itself on its key performance indicator: the relatively low inter-shipment grade variability of its products. The company has been considering alternative production process designs for all stages from mining to ship loading to improve efficiency and reduce costs. Currently product grade variability is well controlled and it is important not to jeopardize the company's reputation for grade reliability. However, the company needs to identify potential efficiencies, reduced costs, and alternative product control systems for potential future expansions. Many process design options from mine face to ship loading must be compared and evaluated, with emphasis on maintaining or improving on current grade variability in final shipments. Pilot studies are infeasible, while complex interactions, competing goals, and numerous system configurations make theoretical analysis unmanageable. Real system trials require strong prior probability of success, particularly where product variability to customers is concerned.

Simulation studies provide a suitable method to evaluate and compare process design alternatives. This paper describes successful simulation modeling of grade variability. The work provides users with easily run Excel models, tailor-made for specific process designs, using Visual Basic macros and graphical output, to explore the implications of decision parameter choice, under a range of possible operating conditions and ore variability. The models have been designed in modular form for the user to interconnect, simulating different configurations of the mine to ship system, to help make informed decisions as to potential system modification. Construction of input mine grade data is a major first step. Meaningful results require input data maintain the correlations, present in the real production environment, between the mineral components, production linkages, and across time. Mining data from real operations are therefore used, with the means and variability adjusted to match potential development proposals. The studies simulate only short-term variability. Long-term variability should not be controlled by the process design but

Material published as part of this publication, either on-line or in print, is copyrighted by the Informing Science Institute. Permission to make digital or paper copy of part or all of these works for personal or classroom use is granted without fee provided that the copies are not made or distributed for profit or commercial advantage AND that copies 1) bear this notice in full and 2) give the full citation on the first page. It is permissible to abstract these works so long as credit is given. To copy in all other cases or to republish or to post on a server or to redistribute to lists requires specific permission and payment of a fee. Contact Publisher@InformingScience.org to request redistribution permission.

rather by longer-term ore extraction plans and ultimately resource availability. It is therefore necessary to filter out medium and long-term variations from the production data before use in the simulation model. A Fourier transform technique to do this is described.

Keywords: Decision Support, Simulation, Information Systems, Industry, Mining.

Introduction

An iron ore mining company in Western Australia has recently established a task force of staff and consultants to refurbish its Product Quality System. Of key importance to the project is improving the efficiency of the production process and reducing costs. In achieving these goals it is essential to maintain the current low variability of shipped product grade. One aspect of the project is to look at alternative process designs from mining to ship loading. The key aims are to achieve low variability of product grade while minimizing costly rehandling of ore by blending onto stockpiles and to develop policies that use knowledge of production rates and logistics to more closely introduce JIT (Just In Time) to each stage of the production line. The criterion for ore quality is to produce ore of consistent grade (not only in iron, but also in phosphorus, silica, and alumina). Traditionally this is achieved by blending from intermediate stockpiles, but this imposes large handling and stock costs which would be avoided if the desired grade uniformity can be achieved with less rehandling of the ore.

The current production process uses large stock inventories and short-term subjective decision making to control product quality by decoupling the mining from the ore handling and crushing process. Whilst highly effective in controlling variability, the system does not take full advantage of the inherent blending that occurs before and after crushing the ore. The system is therefore inefficient because the low variability is achieved by rehandling the ore to an extent that may be unnecessary in a redesigned process. Such rehandling, building, and reclaiming stockpiles through the process from mine to ship is expensive and greatly reduces the profitability of the operation

Improving the process design requires that many options from mine face to ship loading are evaluated and compared. Pilot studies would be prohibitively expensive, while complex mineralogical interactions, competing goals, and numerous possible system configurations limit the applicability of theoretical analysis. It was therefore concluded that simulation modeling would be the most cost effective tool to study how changes to the production process design would affect inter-cargo variability. Everett (1996, 1997, 2001) and Howard, Carson, and Everett (2005) discuss simulation methods used for planning iron ore production projects.

The objective of this paper is to discuss some of the philosophies used in writing the simulation, issues arising during the development, and the solutions implemented to overcome them.

The Flow of Ore and of Information

The company extracts iron ore from three mines. One mine is close to the crusher, and the other two are each several hours trucking distance away. Currently, ore dug up at each mine is put onto stockpiles according to ore classifications. Ore arriving at the crusher from the distant mines is again stockpiled in preparation for blending into the crusher, which reduces the ore particle size to produce two distinct products, lump and fines, each having to closely match their different target compositions, not only in iron but also in phosphorus, silica, and alumina. This quality control is required because the customers' blast furnaces must be fed with ore of consistent composition.

Each day, a crusher feed blend of about 24 kilotonnes is selected from the available stockpiles, to yield crushed products close to target compositions for lump and fines. The products are sampled and assayed after crushing and screening. Crushing alternates between campaigns of standard and low-grade targets, with the campaigns each lasting a week or two. The lump products are re-blended, so three products are exported: PFS, PFL (standard and low-grade fines ore), and PL (lump ore).

The precise composition of the available ore is not known until it is crushed, but decisions based upon estimated compositions have to be made in selecting the ore for crushing. The target com-

positions are for the separate lump and fines crushed products. Samples are taken from the blast holes that are drilled to take the explosive before blasting and mining. These blast hole estimates of the stockpiled ore’s head grade are available when determining the daily blend and are used in combination with bias, lump percent, and lump-fines algorithm models, for estimating the post-crusher product lump and fines grades, as described by Everett, Howard, and Jupp (2010). Crusher grades are available for the lump and fines products as they are produced and sampled. This information is used to assess and revise the status of the overall grade control for generating the next daily blend. An integrated stockpile and blending program (Continuous Stockpile Management System) is used as a decision support tool when generating the blend. The information system can be summarized by the feedback loops of Figure 1.

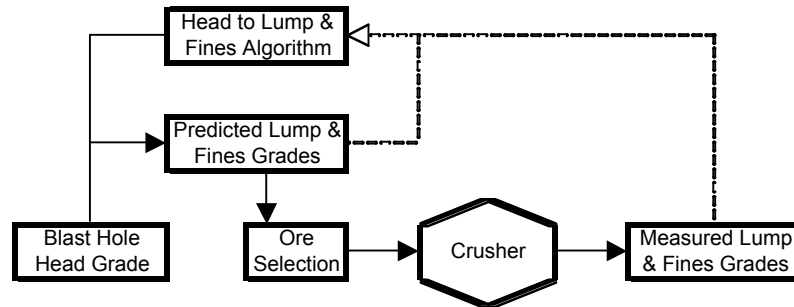


Figure 1. Information flows

Following crushing and screening, the finished product ore is fed onto lump and fines cones. Completed cones are cleared, when possible, directly onto a train, or on to LIFO (Last-In-First-Out) product stockpiles for later train loading.

Finished product is railed a few hundred kilometres to the port, with a schedule of eighteen trains per week. Each train carries only a single product. As much as practical, trains are scheduled to carry product that can be direct loaded when ships are berthed. If not direct loaded to ship, ore is unloaded from the train and stored in one of four sheds. Ore is loaded onto ships by a single ship loader, fed by three front-end loaders or by two, if a train is currently being direct loaded. Shipment cargoes vary in capacity from about 50 to 190 kilotonnes.

No processing of product occurs at the port although significant inherent blending takes place. Decisions on loading ore onto ships are not influenced by product grade. Ore is sampled for

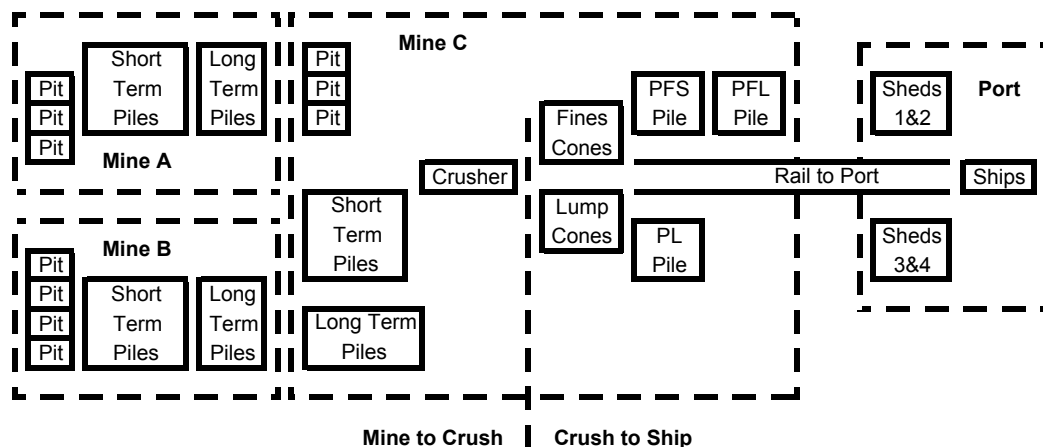


Figure 2: The generalised production system

chemical and physical properties just prior to loading.

Figure 2 shows the flow of ore through the production system, generalized to cover the range of possible process redesigns.

Simulation

A simulation model was designed with the requirement that it be able to simulate not only the current operation, but also minor modifications to the current operation and a variety of alternative operational policies and procedures.

The primary objective was to simulate grade control during normal operation. Accordingly, it was not designed to model breakdowns and interruptions. As a conservative policy, details of operation were ignored, where ignoring them did not overestimate the degree of grade control.

The simulation models were designed, developed, and implemented as Excel workbooks, each with multiple worksheets, backed by Visual Basic (VBA) macros, activated by worksheet buttons. The advantages of using Excel as the programming base lie in the familiarity of the company's personnel with Excel, its excellent graphing capabilities, and its compatibility with VBA. The result is that the simulation models can be run by any of the computer literate stakeholders, with zero marginal cost for extra software licenses and minimal additional training.

Using spreadsheets for serious computer applications is sometimes criticized on the grounds that they can be set up so that data, parameters, and calculations are intertwined, causing confusion and potential loss of integrity. For example, equations can inadvertently be overwritten by numbers, without the operator being aware. To avoid these dangers:

- macros are used principally to carry out calculations instead of using cell equations.
- spreadsheets are protected, except for cells used by the operator to enter parameters.

Both these steps reduce the risk of the user inadvertently corrupting the model. For clarity, the spreadsheets are colour coded to clearly identify the function of the cells to the operator: yellow for editable data, blue for macro-generated results, and white for any other cell function.

Modeling a LIFO stockpile illustrates the convenience of the spreadsheet approach. A portion of a spreadsheet can represent the current occupancy of such a stockpile, as a list with one row per ore record. The most recent input to the stockpile is inserted into the top row of the list, thus moving the rest of the list down. Each extraction from the stockpile is taken from the current top row, unless it exceeds the record's tonnage, in which case it will take the required extra amount from subsequent rows. As rows are exhausted, they are deleted, thus moving the rest of the list up. This method enables the LIFO stockpile to be continuously updated and monitored, much more easily than would be possible by coding a varying dimension array.

Preparing Input Data

To generate input data simulating the ore production, one approach could have been to generate a series of random numbers for each mineral, with appropriate mean and standard deviation (and normalized so that the mineral content totals 100 percent). This approach is inadequate for simulating mining operations, because of the strong cross correlations and serial correlations between the mineral components, between production linkages, and across time.

For the simulation to be effective, it is essential that input data contain the same serial and cross correlations present in the real production environment. Therefore, whenever possible, historical mine production data, from the operations under study, provide the best source of data, because they satisfy this requirement. The use of historical data has to be tempered with the potential

presence of long-term variability in this type of data. In addition, potential changes to mining methods that might affect grade variability and the potential introduction of new pits into the blend must be able to be included in any simulation.

For the simulation, historical data on mining movements on a shift-by-shift basis, over a two-year period, were obtained. These data contained date and shift information, mined tonnage, and blast-block head grade estimates for each movement of ore from each of the pits.

Routine production data are susceptible to input errors so must be tested for validity. Rogue data, lying outside a set multiple of the standard deviation, were identified and real values or best estimates were used in their place.

Production rate

The decoupled production process had allowed mining to be sped up or slowed down independently of daily crushing requirements. The large stockpiles can absorb large variations in the production tonnage. The aim of the simulation was to study a much more uniform production rate, where mining rates were closely aligned to crushing requirements, so it was necessary to adjust the mining data to a more uniform production rate. Long periods of no mining in the data were removed, and the data “stretch” to provide a more constant feed rate. This maintains the advantage of using historical data with their realistic correlations.

Three time based production series of data were generated to gauge the sensitivity of the simulation to the mining delays. Data sets were generated, with delays of greater than 14, 7, or 3 days removed.

The resulting data sets were then manipulated by changing the date and time stamp so that the ore was removed from the pit in the original sequence, but at a more uniform rate over 12 months. This was achieved by adjusting the time (but not the tonnage or grade) of each mined item. To do this, the departure of the cumulated tonnage from a straight line was fitted to the first four harmonic sine waves, and then the timing adjusted to remove these long-period variations in production rate.

Long-term grade trends

Real mining data often contain long-term trends or fluctuations, due to the slow grade changes that occur over the long-term development of a mine. Long-term fluctuations are generally controlled by medium and long-term mine planning.

The production process simulation only reveals effects on short-term variability. Process design and short-term grade control systems cannot control long-term fluctuations unless the ore is stored and rehandled through long-term stockpiles, which the simulation is trying to avoid. If the data used in the simulation contain long-term grade trends, any improvement to short-term variability may be masked by the long-term variability. So, in preparing data for the simulation, long-term fluctuations are filtered out. This is achieved by using the Fourier transform of the grade data, removing the low frequency components, and then recovering the time series using the inverse Fourier transform.

Adjustment of standard deviation and average grade

To simulate changes in mining practices while maintaining the use of historical production data, the standard deviation and the average of the input data could be adjusted. This was done by separating the data into their mean and deviation vectors; adding or subtracting the required amount to the mean vector; multiplying the deviation vector by the required amount, then recombining the

mean and deviation vector to generate a data set having the required means and standard deviations, while maintaining the original correlations between minerals and across time.

By studying a range of variability, the sensitivity of process stockpiling and changes to mining practices can be assessed through their final impact on shipment variability. This permits management to review how trade-offs in cost and effort would affect variability.

For example, a 10% reduction in grade variability coming out of the mine may result in a comparable reduction in shipping variability. This may be made possible, for example, by more frequent digging unit moves and a greater availability of ore sources.

New mine options

The simulation model needs to include mining from possible future pits and, so, requires generation of synthetic data for them. Data are taken from an operating pit most resembling possible outputs from the new mine and manipulated to better represent the proposed pit. Considerations of geology and proposed mining technique are used in the selection process. The average head grade is adjusted to match that established from the resource estimates of the proposed mine, as described above. Production tonnages are to match assumed production rates under life-of-mine plans. The appropriate cross correlations and serial correlations are maintained by this approach.

Varying production rates

The tonnage contribution from each ore source also influences grade variability. Percentage contribution to the daily blend from each contributing source can be varied, as can the total production rate. Similarly, future plans are likely to have different production tonnages from the source data on which they are modeled. So the relative tonnages can be adjusted for of each of the contributing pits, down to the contributions of individual blast blocks.

These methods of preconditioning the data generate a synthetic time-stamped list of floorstocks from each mine. Tonnes and grade are used as input to the mining stages of the simulation models.

The Simulation Modules

The models were designed to simulate variations of policy and procedure relating to production process design, stockpile size, and building strategy. The models were built as two separate modules: “Mine to Crush” output providing the “Crush to Ship” input. This helped testing and modifying the modules as they were written and enables the user to concentrate upon each part of the system, while also being able to run them back-to-back to examine the overall system performance.

Time is used as the common unit for simulation steps (as opposed to tonnage, BCMs, etc.), so the simulation moves forward in steps of a few hours, a shift, or a day as appropriate. All input data and internally generated data simulating activities are time stamped, for example, with the time a stockpile is completed or becomes available. Appropriate movements of ore within the simulation are not initiated until the time stamp permits. The time-based stamps are carried forward within each module and in data transferred between the two modules. The central part of the VBA code is therefore a “do loop” cycling through time units from the start to end of the simulation. This core “do loop” is preceded by code setting up the simulation and followed by code analyzing and reporting the results. For some operations, such as building the train and shipping schedules, the loop repeats iteratively, as it would in designing a real-world schedule.

Key parameters that need to be changed by the operator before a run commences are not coded into the software but are settable in unprotected areas of the control sheets.

The simulations are not designed to fully simulate current production processes, which are quite complex due to short-term non-systematic decision making. Writing the simulation to carry out these current ad hoc procedures would introduce unnecessary complications into the modeling process and yield little increased knowledge of variability.

Tonnage movement is simulated only to the extent necessary to enable grade variability to be understood.

Only steady-state operations are simulated. Interruptions to the production process by such things as significant breakdowns and wet weather are not included. The simulation is not designed to cover special cause events, which are traditionally managed outside the systematic production process.

No emergency stockpiling for significant delays are built into the simulation, but their necessity is acknowledged and will be built into the final process design.

Output from one module to be used as input to the next is transferred automatically, to avoid operator error.

An optional time-stepping function is included, so the simulation can be halted at a particular day and then stepped through a settable range of days to allow study of each day's activity before proceeding to the next day, where it will halt again. This feature is useful not only for debugging the software but also for observing and examining the causes of unexpected results occurring during the simulation.

The simulation is designed to be as simple as possible while not jeopardizing the study of variability. In all cases where a process is simplified, it is ensured that the resulting simulation is conservative and does not underestimate the grade variability. If satisfactory variability can still be obtained, there is no need to introduce extra complexity into the simulation.

The Solver tool in Excel is used to make optimizing decisions in the process that would normally be made by the person responsible for grade control. These decisions are based on the optimization of total grade stress, as discussed by Everett (2007) and Howard and Everett (2008).

There are current long-term stockpiles, which it is desired to eliminate. Accordingly, the simulation allows ore to be fed in from these long-term stockpiles at a fixed but settable percentage of total daily production.

The “Mine to Crush” Simulation Modules

The simulation models were designed so that at the various stages ore may be stockpiled and recovered according to one of three stockpiling disciplines: as a LIFO (Last-in-First-Out), BIBO (Blended-in-Blended-out) or FIFO (First-in-First-out) discipline. Ore mined from the pit is taken sequentially from a list and therefore follows the FIFO discipline. BIBO stockpiles are assumed to be of uniform grade (though with small random variability, as will be discussed below). For this to be achieved they have to be built to completion (lengthwise) and then reclaimed completely (from one end).

Two alternative versions of the “Mine to Crush” simulation module were written to examine the potential of choosing either “Flow through mining” or “Blended stockpiling.”

Flow through mining

This approach aims to maximize the daily amount of ore from each mine going immediately into the daily crusher feed without blending on stockpiles. Ore is mined in a FIFO manner at a rate set by the time stamps of the ore. If ore does not fit the daily blend in the short term (because of unsuitable grade) it is assigned to one of a set of LIFO short-term stockpiles: one for extreme high

grade, one for extreme low grade, and one mid grade. The decision whether ore is to go onto a stockpile is made in the crusher feed management section of the workbook. Solver is used to optimize the daily crusher blend, across all of the hubs, by choosing from all the extracted ore produced and in short term stockpiles on that day. A settable penalty ensures most of the mid grade ore should flow through the process and not onto short-term stockpiles. If mined ore is sent to a LIFO pile, ore from another pile will replace it in the production pipeline unless the ore was excess to daily blend tonnage requirements.

To set the grade ranges for the LIFO piles, the percentage of ore to be classified as extreme grade, high or low, is set before the simulation is run. For each hub, the first principal component (PC1) of the grade vector (comprising all the analytes making up the ore grade) is extracted from all of the pit data for each hub. PC1 is the linear composite which best explains the variability of all the minerals, maximizing the eigen-value, defined as the sum of the squared correlations (or “loadings”) with the mineral grades. Each of the historical movements of ore from each pit is given a PC1 percentile score (0% to 100%). Ore scoring above, between, or below settable percentile limits goes to the high, mid, or low-grade LIFO pile, if it is not immediately hauled to crusher. The percentage of ore going to the extreme stockpiles is thus easily adjusted by changing the settable limits.

Blended stockpiling

This mining and stockpiling method extracts all ore from the mine as it is produced, in time stamp sequence, and stacks it onto paired BIBO stockpiles (one being stacked while the other is reclaimed), located at each of the hubs. Once a stockpile reaches its nominated tonnes it is closed off and is then available to be reclaimed until empty, while a new stockpile for the same grade range begins.

In the simulation control sheets, the number of stockpile pairs, their grade range, and tonnage capacity can all be set. Separation of ore into the stockpile pair grade ranges is based on settable cut-offs, using the blast head grade estimates. A pre-run shows the percentage of ore that will go through each BIBO pair of stockpiles. Once completed, the stockpile carries the weighted average grade of the ore that has been added to the pile and the time stamp of its completion. The crusher feed management module uses the blended stockpile grade, tones, and time stamp from each hub to create the daily crusher blend, again using the optimization function of Solver.

Using Solver for crusher feed management

The simulation reads in mined ore from worksheets containing preconditioned historical data and passes it on to the crusher feed management section, or via short-term piles, as required by the mining concept being simulated. The model then uses the Excel facility “Solver” to select each day’s crusher feed, to give an exponentially smoothed forecast grade as close as possible to the target (where the target is automatically the average grade for the whole simulation period, since long-term trends have been removed). Solver makes ore grade selections to minimize total grade stress as discussed by Everett (2007) and Howard and Everett (2008). The operator sets parameters, such as the daily production from each mine, the quantity of ore to be trucked, and the percentage of ore from long-term stockpiles, which act as constraints when running Solver. For each day’s crusher feed, Solver carries out multiple calculations until it reaches an optimum blend of floorstocks and or stockpiles, satisfying the constraints and moving the exponentially smoothed grade as close as possible to target for both lump and fines.

If ore is not available from the pit, because of a delay in mining, the ore to meet production requirements will be automatically taken from the LIFO piles. If no ore is left in the LIFO piles, the simulation will offer the option of taking ore from the long-term stockpiles that are available to meet the tonnage target. The amount of ore taken from the long-term stockpiles is recorded and

graphed. Random grade variability can be introduced for each long-term stockpile output to better simulate real outcomes of reclaiming ore from stockpiles of this type.

The daily crusher product tonnages are output, providing input to the following “Crush to Ship” module.

The “Crush to Ship” Simulation Module

The “Crush to Ship” module simulates the flow from crusher output, through railing, port storage, and ship loading, tracking the tonnage and grade history along the way. Output from either of the “Mine to Crush” modules provides input to this module. No grade decisions are made in this module, as the procedure is generally fixed, so Solver is not required. The inter-cargo variability for each product, and for each mineral, is reported as a standard deviation and as a “Process Capability” (defined as the target standard deviation divided by the achieved standard deviation).

This module is set up to perform three broad functions, in the following sequence:

1. Generation of the shipping and railing schedule

In all stages of the simulation, tonnages must be matched to avoid large build-ups or deficits of ore. A shipping and railing schedule is generated, matching the tonnages produced through the crusher. To avoid tonnage build-ups, ship arrivals are scheduled to match production rates and to take the product with most stocks. The percentages can be set for ships of two sizes (Panamax, 85 kt capacity and Cape, 160 kt capacity). If a large ship is to take two cargos, the product for each cargo is independently determined, based on the above criterion. This process is designed to mimic the development of the annual shipping plan and the quarterly nomination system in its simplest form. As in reality, the process is iterative: first setting a train schedule, then a shipping schedule, then revising the train schedule to determine the specific products to be railed.

The schedule of train arrivals at the crusher each day is settable, with a random component to mimic variability in achieving the schedule. The product for each scheduled train is selected on two criteria that mimic the real scheduling process in a simplified form:

- a) If arriving at the port while a ship is berthed, carry the product for that ship, and,
- b) If there will not be a ship at berth, carry the product of highest tonnage at the crusher. Because tonnages are matched throughout the simulation this is equivalent to railing based on the product tonnages most required at the port.

2. Building post crusher stockpiles and loading trains

The “Load Train” stage of the module feeds the product ore from the previous module through the cones, loads the trains, computes their grades and tracks the cone and post-crusher stockpile tonnage and grade history. As far as possible, the simulation loads the trains directly from the crusher output cones. If the dead cone ore is not of the required product to be railed, it is fed into LIFO product stockpiles in a timely manner to ensure the cone footprint is empty when the live cone is full. Ore is also loaded to trains from product stockpiles to obtain the loading rates required. Ore from cones has an error applied as will be described in the “Random errors and events” section below.

Figure 3 shows an example plot of the history for a post-crusher stockpile.

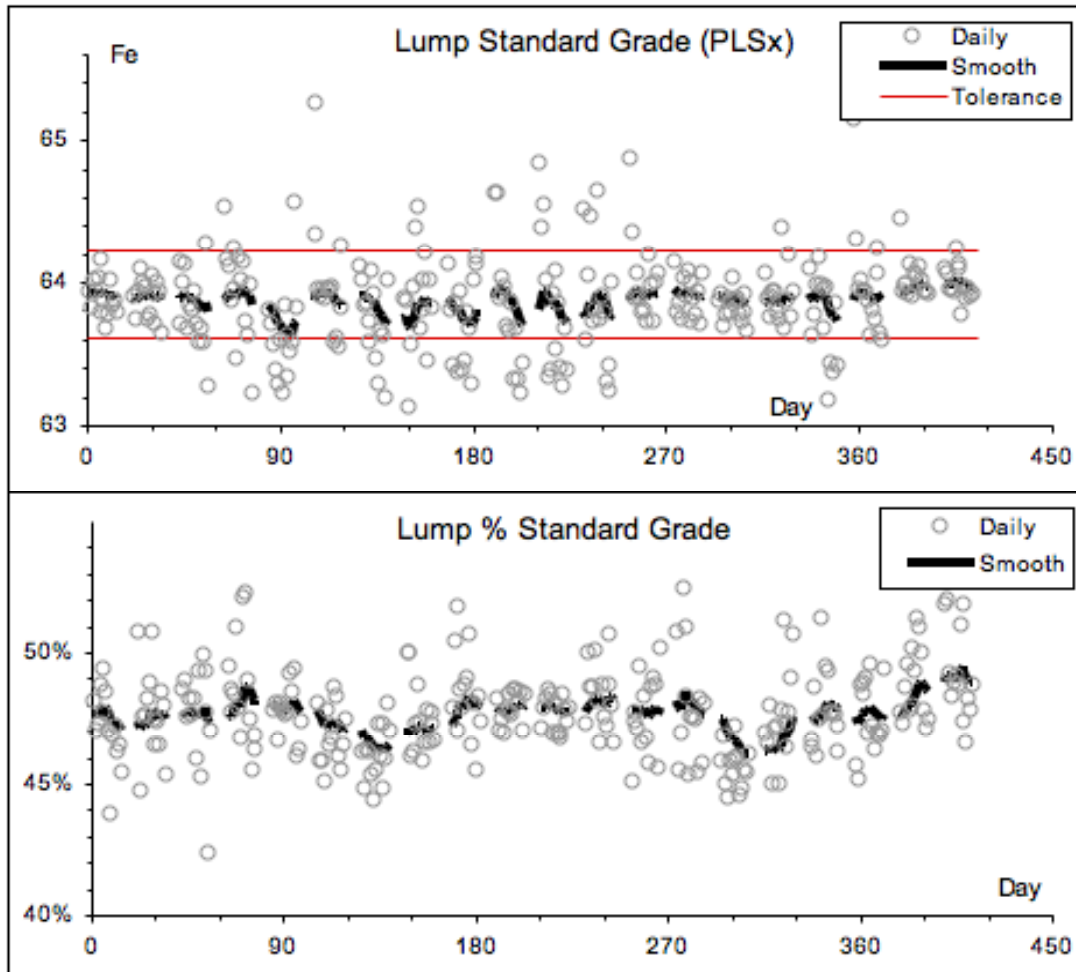


Figure 3: Post-crusher stockpile history

3. Train to ship

Trains arriving at the port are dumped and loaded direct to ship if a ship is berthed. If there is no ship, they are stacked to one of the four sheds, subject to constraints designed to ensure any product is available for appropriate loading from both pairs of sheds. Stacking and reclaiming of ore from the sheds follow a LIFO discipline within the simulation. This is a simplification but conservative in its effect on grade variability. Ships arrive and are loaded in accord with the shipping schedule, at a constant settable rate, matching the average rates achieved in practice. The rates differ for loading direct to ship or from the sheds.

Setting Key Parameters in the Simulation Process

For each simulation run, several key parameters need to be input, including stockpile sizes and disciplines, daily production rates, the proportion of long-term stockpile ore to be included, and the target grades and tolerances. They are entered on the simulation module's "Control" sheet.

Random inputs

Random number sequence seeds are set for where the simulation models require random numbers (as in the generation of the train and shipping schedule). If these seeds are changed, a different random number sequence will be generated for the next simulation, but otherwise successive simulation runs are identical. For debugging, it helps to be able to rerun the same random sequence.

In recovering ore from blended stockpiles, as in the case of the BIBO piles at the mine and the crusher cones, using the average grade will underestimate variability. A settable random variability (based upon a study of real data) is therefore added into the data as such stockpiles are reclaimed.

Running the Simulation

Since the output from one provides the input to the next, each simulation module can be run separately. Sections of each module can also be run independently.

A linking worksheet assists the user to automatically run a scenario from mine to ship. A range of options, such as varying the contribution of long-term stockpiles or the size and grade of stockpiles, can be automatically completed with key summary reports generated for each run. This way, a fairly lengthy set of runs (requiring several hours of computation) can be achieved without potential operator intervention and error.

Parameters and simulation results for each module are concisely reported in a summary sheet, which can be stored after each run and compared to results from subsequent runs. Important key performance indicators are graphed within each module. Figures 4 and 5 show examples from the “Mine to Crush” and “Crush to Ship” modules. All data and outputs are stored for each run, enabling operators to construct different sets of graphs or extract the data to another workbook for further study.




Figure 4: Example crusher product history

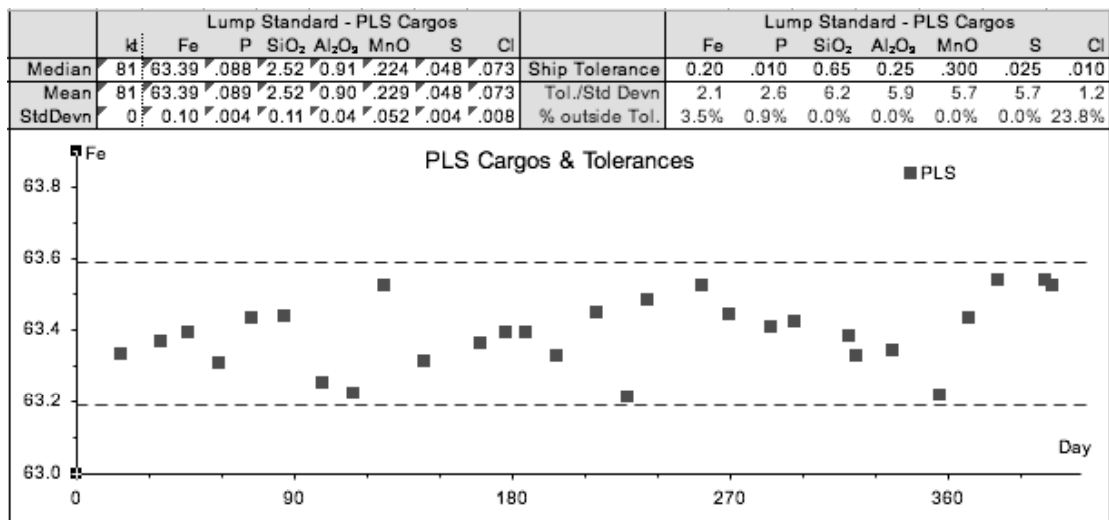


Figure 5: Shipping grade report




Figure 6: Shed occupancy at the port

Figure 6 graphs the history of the tonnages for two products stored in one of the port sheds. The tonnage for one of the products is plotted downwards, so the gap between the two traces corresponds to the unused shed capacity. The assumed shed capacity is 300 kt.

The simulation modules are designed to be understandable and useful to a variety of operators. Each module contains a summary sheet of key results from any run and graphs of key performance indicators. The outputs cover a range of levels, spanning, for example, from detailed stockpile histories for the operators to overall key performance indicators (such as shipping grade variability) for management.

Conclusion

The use of real production data maintains serial and cross correlations between the mineral components, production linkages, and across time. Adjusting the standard deviations and means of the data as well as creating hybrid data sets, with delays removed, creates multiple scenarios that can be modeled for sensitivity analysis.

The simulation modules described have the capability of modeling potential new process designs for the company's operations and to gauge the effect on inter-shipment variability. They allow

management to make policy decisions for revising the process, taking into account cost implications and the effect on grade variability of ship cargoes, a the vital key performance indicator.

The model has simplified the process. Simulation of grade performance rather than tonnage performance has been borne in mind throughout. Where there have been conscious simplifications, these have been done in a conservative manner with respect to reported grade variability, thereby not overestimating the achievable grade control.

The results of the simulation were accepted by the company and have enabled some considerable cost saving operational changes that otherwise would have been very difficult to justify quantitatively. To test the robustness of the simulation models, realistic scenarios were run, in cooperation with experienced mining engineers and geologists. The changes justified by simulation have paid off on implementation.

One great benefit of using VBA behind a spreadsheet workbook is that the mining engineers and geologists are familiar with Excel. They can run scenarios themselves in a familiar environment (and without licensing problems). This factor greatly helped the acceptance of the study by management and operational staff.

The simulation methodologies discussed are readily applicable to other mining and processing operations.

References

- Everett, J. E. (1996). Iron ore handling procedures enhance export quality. *Interfaces*, 26(6), 82-94.
- Everett, J. E. (1997). Simulation to reduce variability in iron ore stockpiles. *Mathematics and Computers in Simulation*, 43, 563-568.
- Everett, J. E. (2001). Iron ore production scheduling to improve product quality. *European Journal of Operational Research*, 129, 355-361.
- Everett, J. E. (2007). Computer aids for production systems management in iron ore mining. *International Journal of Production Economics*, 110/1, 213-223.
- Everett, J. E., Howard, T. J., & Jupp, K. (2010). Combined regression and exponential smoothing methods to predict crusher grades and lump percentage. *Mining Technology*, 118/2, 102-108.
- Howard, T. J., Carson, M. W. & Everett, J. E. (2005). Simulation modeling of iron ore product quality for process and infrastructure development. In *Proceedings MODSIM 2005 International Congress on Modelling and Simulation*, 1251-1257. Modelling and Simulation Society of Australia and New Zealand: Canberra.
- Howard, T. J., & Everett, J. E. (2008). Maintaining product grade from diverse mine sites at BHP Billiton Iron Ore Newman Joint Venture. *Mining Technology*, 117/1, 12-18.

Biography



Jim Everett began as a research geophysicist at Cambridge and the Australian National University. His family always said “what are you going to do when you grow up,” so he joined industry, as a petroleum exploration geophysicist. With a computer background, he was put onto project evaluations. Losing arguments with accountants, he did a couple of Economics and Commerce degrees at the University of Western Australia (and found the accountants were usually right). The university was starting its MBA course and wanted people with industry experience to teach on it. With a young family, fieldwork was no longer so attractive, so Jim jumped at the chance of returning to academia. Thirty years on, he has now retired as Emeritus Professor of Information Management and divides his time between research and consulting to Western Australia’s booming mining industry.