Optimised Cell Batching for New Zealand Aluminium Smelters Ltd

David M Ryan
Department of Engineering Science
University of Auckland
New Zealand
d.ryan@auckland.ac.nz

Abstract

New Zealand Aluminium Smelters Ltd operates a smelting facility at Tiwai Point near Invercargill. The smelter produces aluminium by the electrolytic reduction of alumina in 650 cells. The cells are laid out in long lines and are grouped into tapping bays. The purity of the aluminium varies from cell to cell depending on a number of factors including the age of the cell, the purity of the alumina feed and the manner in which the cell has been operated during its production life. High purity aluminium commands a premium price on the metals market. In the production process molten aluminium is tapped from the cells into crucibles that are used to transport the metal to furnaces from where finished product is cast in the form of billet or ingot. Each crucible taps the metal from three cells that are located within some specified limit of spread in the same a tapping bay. The spread limit is imposed to minimise the time taken to fill the crucible. This paper describes the development of a set partitioning optimisation model to batch or group triples of cells so that the total value of metal produced is maximised. The main aim is to minimise the dilution of high purity (high value) metal by low purity (low value) metal. The optimised batches tend to group high purity cells together and leave the lower purity cells to be batched with other lower purity cells. Numerical results show that significant improvements in excess of 10% can be achieved in the value of metal by carefully batching cells.

1 Introduction

New Zealand Aluminium Smelters Ltd operates a smelting facility at Tiwai Point near Invercargill. The smelter produces aluminium by the electrolytic reduction of alumina according to the reduction equation $2Al_2O_3 + 3C \rightarrow 4Al + 3CO_2$. This reaction which is called the Heroult-Hall process, is carried out in reduction cells constructed of an outer steel shell and a lining of refractory bricks. A carbon cathode is placed in the floor of the cell and carbon anodes are suspended above the cell on cast iron yokes. A very high direct current of approximately 190,000 amps is passed between the anode and cathode through a bath of molten cryolite at 960^0C which provides the electrical conductivity. The alumina is feed into the cryolite bath at regular intervals from a hopper that is located above the cell. The carbon required in the reduction reaction is provided by the carbon anode blocks which gradually reduce in size over a period of approximately twenty-seven days. When a block becomes too small, it is replaced by a new block. As the aluminium is produced it sinks to the bottom of the cell. Each day approximately 1260kgs of molten aluminium are tapped from the cell into a crucible by a vacuum siphoning system. A crucible is a large steel bucket lined with refractory bricks. Each crucible can tap the aluminium from three cells.

The cells are laid out in four lines. Three of the lines are each approximately 600 metres long and consist of 204 cells grouped into four tapping bays each made up of 51 cells. The fourth line 300 metres long was installed more recently and is made up of 48 cells of a newer technology in one tapping bay. All the cells in a tapping bay are tapped once each day and produce seventeen (or sixteen in the case of line 4) crucibles. Each bay is tapped either during the day shift or during the night shift. Once each crucible is filled with aluminium from three cells (always from the same tapping bay), it is transported from the reduction lines to furnaces in the Metal Products Division from where it is cast into finished products in the form of ingot or billet.

The purity of the aluminium varies from cell to cell depending on a number of factors including the age of the cell, the purity of the alumina feed and the manner in which the cell has been operated during its production life. The purity of the aluminium declines gradually as contaminants in the form of iron, silicon, gallium and other chemicals increase until at some stage a decision is made to cease production in the cell. The cell is then taken off-line and rebuilt before being brought back into production some days later. Each cell is assayed regularly to determine the percentage of aluminium, iron, silicon, gallium and other chemicals. Because high purity aluminum commands a premium price on the metals market, it is important that aluminium tapped into a crucible from high purity cells is not contaminated by tapping from low purity cells into the same crucible.

This paper describes the development of a set partitioning optimisation model to batch or group triples of cells so that the total value of metal produced is maximised. The main aim is to minimise the dilution of high purity (high value) metal by low purity (low value) metal. The optimised batches tend to group high purity cells together and leave the lower purity cells to be batched with other lower purity cells. Numerical results show that significant improvements in excess of 10% can be achieved in the value of metal by carefully batching cells.

In Section 2 of this paper, we will describe an optimisation model for cell batching and discuss the formulation of a natural objective to measure the solution quality. In Section 3, we will discuss aspects of the solution process and in particular outline how integer solutions can be derived from continuous LP relaxation solutions. Some numerical results will be presented in Section 4 and in Section 5, we will briefly outline the relevance of the cell batching model in the context of a wider furnace scheduling problem which is the subject of on-going research and development.

2 An Optimisation Model for Cell Batching

The cell batching optimisation can be formulated naturally as a set partitioning problem (SPP) which can be written as

$$\text{minimise } z = \mathbf{c}^T\mathbf{x}, \quad A\mathbf{x} = \mathbf{e}, \quad x_j = 0 \text{ or } 1$$

where A is a 0-1 matrix and $\mathbf{e}^T = (1,1,...,1)$. Because cells in different tapping bays can never be tapped into the same crucible, the cell batching optimisation problem for each tapping bay can be considered independently of the other bays. For each tapping bay, the 51 constraints or rows of A correspond to cells in the tapping bay and ensure that each cell appears in exactly one batch or crucible. The columns of A represent all possible triples of cells (i.e. batches) which could be tapped into the same crucible. Each column then has exactly three nonzero unit values. For example, a batch made up of cells 1, 3 and 6 would be represented in the model by a column with zeros everywhere except for unit values in rows 1, 3 and 6. In general then the elements a_{ij} are defined as

$$a_{ij} \begin{array}{l} = 1 \text{ if cell i is included in batch j and} \\ = 0 \text{ otherwise.} \end{array}$$

In this basic unrestricted form of the model there are ${}^{51}C_3$ or 20825 columns or variables which can be easily enumerated. A solution of this SPP will be made up of exactly seventeen variables at unit value (representing the chosen batches) and all other variables will have zero value.

2.1 Spread Limitations on Cell Batches

Because each tapping bay is approximately 150 metres long, it is not practical to tap cells that are far apart into the same crucible. The actual tapping process involves the use of a gantry crane to carry the crucible and the human operator walks along the line from cell to cell. To avoid requiring the operator to walk large distances to fill each crucible, the usual practice is to tap cells that are within some specified maximum distance apart on the line. This is referred to as the spread of the batch. Spread can be defined simply as the difference between the maximum and minimum cell number in the batch. So three adjacent cells have a minimal spread of two. A batch made up of cells 1, 3 and 6 would have a spread of 5. In the enumeration of the columns or batches in the SPP model, the spread for each batch can be calculated. If the spread exceeds a specified limit, the batch can be rejected and not included in the model. Alternatively, the batch could be included in the model but marked as having a spread exceeding the maximum spread. During the optimisation process, these batches would be ignored unless the maximum spread was increased sufficiently. Such an increase could easily be included in a post-optimal investigation. In Section 3 we comment further on such investigations.

Batches exceeding the maximum spread can also be considered in a more useful manner. By adding an additional generalised SPP constraint to the basic SPP model, we could permit a limited number of batches with excessive spread to be included in the solution. This reflects the management view that a small number of batches (e.g. one or two of the seventeen batches) with excessive spread can be tapped provided they generate a sufficiently improved optimal solution when compared to the optimal solution using only batches that are within the maximum spread limit. Typical spread limits might be a maximum spread of 5 but up to two batches with a maximum spread of 10. In the enumeration of columns in the SPP model, all batches with a spread up to 10 would be generated but all those batches with spread between 6 and 10 would contribute to the additional constraint with a right-hand-side of 2. These spread restrictions which are included either implicitly (i.e. no spread exceeding 10) or explicitly (i.e. limited spread exceeding 6) significantly reduce the total number of variables in the SPP to something less than about 3000.

2.2 Alloy Codes and a Cell Batching Objective Function

A natural objective for the cell batching optimisation can be based on some estimate of the market value of the aluminium. This will obviously reflect the purity of each batch. Batch purity can be calculated as a simple weighted average of the known cell purities which make up the batch where the weights reflect the weight of metal tapped from each cell. Although in practice the actual cell tapping weights do vary a little, it is reasonable to assume before the tapping takes place that the tapping weights will be constant at 1260kgs for lines 1 to 3 and 1480kgs for line 4. Given constant tapping weights and the cell assay values for aluminium, iron, silicon and gallium, the batch values are calculated as simple averages.

Table 1 defines eighteen aluminium alloy codes with their corresponding minimum aluminium percentage and maximum percentages for silicon, iron and gallium and an estimate of the corresponding alloy premium value. Generally speaking as the aluminium percentage increases (with corresponding decreases in the silicon, iron and gallium percentages) the alloy premium value increases rapidly. While this Table refers more particularly to finished product values cast in the particular alloy codes, we can classify the batch purity using the alloy specifications and then use the premium value as an objective coefficient for the batch. The negative premium values for the first three alloy codes reflect the fact that these grades of aluminium usually require purification by mixing with a higher grade metal and thus result in an actual loss of premium.

Code	Min Al%	Max Si%	Max Fe%	Max Ga%	Premium
AA????	0.000	1.000	1.000	1.000	-50.00
AA150	99.500	0.100	0.300	0.100	-40.00
AA160	99.600	0.100	0.300	0.100	-25.00
AA1709	99.700	0.100	0.200	0.100	0.00
AA601E	99.700	0.100	0.080	0.100	40.00
AA601G	99.700	0.100	0.080	0.100	40.00
AA185G	99.850	0.054	0.094	0.014	15.00
AA190A	99.900	0.054	0.074	0.014	45.00
AA190B	99.900	0.050	0.050	0.014	50.00
AA190C	99.900	0.035	0.037	0.012	110.00
AA190K	99.900	0.045	0.055	0.034	100.00
AA191P	99.910	0.030	0.045	0.010	120.00
AA191B	99.910	0.030	0.027	0.012	139.00
AA192A	99.920	0.030	0.040	0.012	140.00
AA194A	99.940	0.020	0.040	0.007	150.00
AA194B	99.940	0.034	0.034	0.010	180.00
AA194C	99.940	0.022	0.027	0.009	200.00
AA196A	99.960	0.020	0.015	0.010	260.00

Table 1. Alloy Code Specifications and Premiums.

2.3 Off-line Cells

When cells are taken off-line to be rebuilt, they will not be tapped. This implies that one or more of the seventeen batches from the tapping bay will include fewer than three cells. It is important that the optimised solution determine which batches should be composed of fewer cells. One approach would be to generate all possible batches involving one and two cells and include them in the SPP model when the tapping bay has off-line cells. This results in a very large increase in the number of variables and causes further computational problems during the solution process. A much more attractive approach is to simply treat off-line cells as having a zero tapping weight. The off-line cells are then permitted to appear anywhere in the cell order during the enumeration of batches. In other words, the off-line cell can appear in any triple of cells without affecting either the spread calculation or the batch chemical composition. For example, if cell 50 is off-line, then a batch made up of cells 1, 2 and 50 would have a spread of 1 and the chemical composition would be determined entirely by cells 1 and 2. If this batch were included in the optimal solution, it would be interpreted as a batch involving just cells 1 and 2. However the SPP constraint for cell 50 would have been satisfied by this variable. The advantage of this approach is that all batches remain triples of cells including all triples involving off-line cells and the SPP model (at least for lines 1 to 3) is always made up with 51 cell constraints.

3 The Solution Process

The SPP model is solved by first solving the LP relaxation problem in which the integer restrictions are relaxed. The LP solution is trivial. With little extra effort it is possible to report on a sequence of LP solutions which gradually include variables with wider and wider spreads. During this initial optimisation phase, batches with spreads exceeding the maximum permitted spread are ignored or equivalently, the right-hand-side for the additional constraint described in Section 2.1 is set to zero thus preventing batches with excessive spread from contributing to the solution.

3.1 Handling Excessive Spread Batches

It is also simple to quantify the benefits of permitting a small number of batches with spreads exceeding the specified maximum spread as discussed in Section 2.1. After completing the initial LP solution, the right-hand-side for the additional constraint can be increased slowly to identify the potential benefits of using a limited number of batches with wider spread. All of these calculations are performed using the LP relaxation. While the LP solutions do exhibit some evidence of natural integer structure, most solutions involve fractional variables that must be forced to integer values using a branch and bound algorithm.

3.2 Constraint Branching in the Cell Batching Optimisation Model

The natural constraint branch (see Ryan and Foster (1981)) for this SPP is defined by any pair of cells that are not always included together in batches in the fractional solution. It is easy to show using balanced matrix theory (see Ryan and Falkner (1988)) that such a pair of cells must always exist in any fractional solution. The binary constraint branch can then be imposed by requiring on the one-side of the branch that the two cells always appear together in a batch and on the zero-side of the branch that the two cells always appear in different batches. The implementation of the branch involves removing sets of variables from the descendant LPs. On the one-side, batches in which the cells do not appear together are removed (effectively by setting their corresponding variable upper bounds to zero) and on the zero-side batches in which the cells appear together are removed.

While this constraint branching strategy is particularly effective, it is true that after imposing a sequence of constraint branches, the LP can become infeasible because a cell becomes isolated from its neighbours in such a way that no feasible set of batches can include the cell. When this happens it usually results in a long sequence of fathoming infeasible nodes and the branch and bound process can take a long time to find a feasible integer solution. For this reason we have implemented a heuristic integer allocation process at each node in the branch and bound tree in order to find integer solutions more quickly.

3.2 Integer Allocation Heuristics

Because of the special structure of the underlying SPP model in this application, it is easy to create infeasible LPs at nodes in the branch and bound tree. To avoid the computational problems that this causes, we have implemented an integer allocation process which is applied at each fractional node in the branch and bound tree including the root node. The process is heuristic in that it attempts to force an integer solution from the fractional solution by making a sequence of greedy decisions. First, all variables (batches) at value one in the fractional solution are fixed at that value. We then find the first cell which is not yet included in a batch and search amongst all variables (including those which are nonbasic) for the least cost batch including that cell and two other cells which are also yet to be covered. If no such variable can be found, the search is abandoned and the integer allocation fails. The allocation process can be applied using any order of the cells. In our implementation we apply the allocation process considering the cells in increasing and also in decreasing order.

The process is remarkably effective in that it often produces integer solutions with objective values very close to the LP upper bound value. Such an integer solution often becomes the best bound that fathoms the remaining live nodes in the branch and bound tree and the solution process terminates.

4 Some Numerical Results

We report here some results that illustrate the performance of the cell batching optimisation and compare the results with a so-called default order solution. The default order solutions are based on batching cells in the natural sequence such as (1,2,3), (4,5,6), (7,8,9) etc. The default order solutions actually form the basis of the manual solution method in which obvious poor value batches are modified locally by changing the allocations of cells to nearby batches to improve the solution. The following Tables include cost values based on the alloy code premium for each batch. It can be seen that in tapping bay 2AW, the optimised solution produces 6 batches with premium value of 40.0 (alloy code AA601E) while the default solution produces just one batch with this premium value. Notice also that batches 12 and 14 in the optimised solution involve spreads of 5 and 7 respectively. All other batches have spreads not exceeding 4. The two off-line cells (366 and 392) appear in batches 6 and 9 respectively as negative cell numbers. The weights, spreads and chemical compositions of these two batches ignore the off-line cells.

Optimised solution for tapping bay 2AW: [Cells 352 to 402] [Dataset d271197]												
Cell 15(-366) is off-line												
Cell 41(-392) is off-line												
A total of 51 cells found for bay 2AW; 2 cells off-line												
Generating with MINSPREAD 4; MAXSPREAD 10; MAXSPREADRHS 2												
Model size: 52 constraints; 2237 variables												
Bay	#	Weight	%AL	%SI	%FE	%GA	Code	Spread	Prem	Cells		
2AW	1	3840	99.800	0.044	0.095	0.016	AA1709	2	0	352	353	354
2AW	2	3840	99.810	0.040	0.105	0.016	AA1709	2	0	355	356	357
2AW	3	3840	99.737	0.041	0.167	0.015	AA1709	4	0	358	360	362
2AW	4	3840	99.847	0.042	0.080	0.014	AA601E	4	40	359	361	363
2AW	5	3840	99.760	0.041	0.132	0.016	AA1709	3	0	364	365	367
2AW	6	2560	99.835	0.034	0.079	0.015	AA601E	3	40	-366	399	402
2AW	7	3840	99.813	0.036	0.101	0.016	AA1709	2	0	368	369	370
2AW	8	3840	99.823	0.038	0.096	0.015	AA1709	3	0	371	372	374
2AW	9	2560	99.840	0.034	0.080	0.015	AA601E	2	40	373	375	-392
2AW	10	3840	99.820	0.038	0.095	0.016	AA1709	2	0	376	377	378
2AW	11	3840	99.833	0.035	0.087	0.016	AA1709	2	0	379	380	381
2AW	12	3840	99.837	0.037	0.079	0.015	AA601E	5	40	382	384	387
2AW	13	3840	99.817	0.037	0.093	0.016	AA1709	3	0	383	385	386
2AW	14	3840	99.837	0.033	0.080	0.016	AA601E	7	40	388	389	395
2AW	15	3840	99.753	0.075	0.126	0.015	AA1709	3	0	390	391	393
2AW	16	3840	99.843	0.037	0.080	0.015	AA601E	3	40	394	396	397
2AW	17	3840	99.827	0.036	0.092	0.016	AA1709	3	0	398	400	401
Integer objective: 240 (Branch and Bound time of 12.82 seconds)												
Default 2 spread solution for tapping bay 2AW: [Cells 352 to 402]												
2AW	1	3840	99.800	0.044	0.095	0.016	AA1709	2	0	352	353	354
2AW	2	3840	99.810	0.040	0.105	0.016	AA1709	2	0	355	356	357
2AW	3	3840	99.770	0.041	0.153	0.013	AA1709	2	0	358	359	360
2AW	4	3840	99.813	0.041	0.094	0.016	AA1709	2	0	361	362	363
2AW	5	2560	99.760	0.040	0.121	0.016	AA1709	1	0	364	365	-366
2AW	6	3840	99.793	0.038	0.119	0.016	AA1709	2	0	367	368	369
2AW	7	3840	99.827	0.035	0.093	0.015	AA1709	2	0	370	371	372
2AW	8	3840	99.830	0.038	0.090	0.015	AA1709	2	0	373	374	375
2AW	9	3840	99.820	0.038	0.095	0.016	AA1709	2	0	376	377	378
2AW	10	3840	99.833	0.035	0.087	0.016	AA1709	2	0	379	380	381
2AW	11	3840	99.820	0.040	0.090	0.016	AA1709	2	0	382	383	384
2AW	12	3840	99.833	0.034	0.083	0.016	AA1709	2	0	385	386	387
2AW	13	3840	99.820	0.035	0.087	0.016	AA1709	2	0	388	389	390
2AW	14	2560	99.715	0.093	0.144	0.015	AA1709	2	0	391	-392	393
2AW	15	3840	99.847	0.037	0.077	0.015	AA601E	2	40	394	395	396
2AW	16	3840	99.837	0.035	0.083	0.015	AA1709	2	0	397	398	399
2AW	17	3840	99.837	0.035	0.088	0.016	AA1709	2	0	400	401	402
Default 2-spread objective: 40												

Table 2: Optimised and Default Solutions for Tapping Bay 2AW

In Table 3 we give a comparison of the overall results of the optimised cell batching applied to all tapping bays. The optimised results can be compared with the corresponding overall results of the default solution. The results are reported in terms of the number of batches (i.e. crucibles) produced in each alloy code and the total premiums generated by those batches. The optimised solutions show an improvement of approximately 11% over the default solutions.

		Optimised solution		Default solution	
Code	Premium	# crucibles	premium	# crucibles	premium
AA????	-50.0	2	-100	2	-100
AA160	-25.0	5	-125	7.67	-191.67
AA1709	0.0	114	0	122.33	0
AA601E	40.0	27	1080	18	720
AA185G	15.0	2	30	2	30
AA190C	110.0	1	110	2	220
AA190K	100.0	4	400	4	400
AA191P	120.0	4	480	4	480
AA192A	140.0	10	1400	14	1960
AA194B	180.0	21	3780	16.33	2940
AA194C	200.0	12	2400	13.67	2733.33
AA196A	260.0	14	3640	10	2600
Totals:		216	13095	216	11791.7
Percentage increase in total default premiums 11.05%					

Table 3. Summary by Alloy Codes for all Tapping Bays

5 Conclusions

While the cell batching optimisation described in this paper has produced significant improvements in the value of metal produced from the reduction lines, this problem really forms part of a larger production scheduling problem at the smelter. The full problem involves first a decision about which products from the order book to produce during the day. This production must then be scheduled on the furnaces and casting machines in the Metal Products Division. A furnace production schedule is made up of periods during which the furnace is filled with suitable batches from the reduction lines followed by periods during which the chemical composition of the furnace metal is adjusted and stabilised before the required product is cast. Given a production schedule for each furnace, the allocation of batches from the cell batching optimisation to furnace fills must be decided. This particular problem is not trivial since it implies a time sequencing for the actual production of the batches. There are particularly important constraints on the production of batches that result from the limited capacities of the gantry cranes used during the tapping process. In the cell batching optimisation, these constraints were ignored as was the time sequencing of production of the batches. We are currently investigating this further scheduling problem to determine a feasible production sequence for the optimal batches to match a given furnace schedule.

Acknowledgment

The work reported in this paper was begun by Tuck (1997) in his Year Four Project in Engineering Science.

References

[1] D.M. Ryan, J.C. Falkner, *On the Integer properties of Scheduling Set Partitioning Models*, European Journal of Operations Research, 35 (1988), pp442-456.

[2] D.M. Ryan, B.A. Foster, *An Integer Programming Approach to Scheduling*, in A. Wren (ed) Computer Scheduling of Public Transport, North Holland, (1981), pp 269-280.

[1] S. Tuck, *Optimal Cell Batching for New Zealand Aluminium Smelters Tiwai Point Facility*, Year Four Project, Dept of Engineering Science, University of Auckland, (1997).