COPPER ORE HEAP LEACHING

N.G. Barton[1] and G.C. Wake[2]

Heap leaching of ore is becoming more important as a cost effective means of metal extraction. Important metals like copper and gold are produced by this process. In this report, we study chemical and physical processes in heap leaching in order to gain a better understanding of leach heaps. These processes are complicated, not well supported by data, and generally not well understood. Heap leaching is investigated using a variety of complementary approaches: data analysis, microstructure models of the ore, overall chemical processes, and analysis of fluid motions by continuum and probabilistic methods. A significant finding of our work is that voids in the heap are not full of solution, and further the distribution of leach liquor within the ore is not uniform which decreases the efficiency of the leaching process. We suggest mechanisms by which the uniformity of the leach liquor distribution can be increased, thus increasing the efficiency of the process.

1. Introduction

WMC (Copper Uranium Division) asked the Study Group to investigate the performance of their heap leaching operation at Nifty, Western Australia, approximately 400 km SE of Port Hedland. This plant was designed to produce 16.5 Kt of copper per year. The copper leach heap in the pilot study is 120 metres long, 45 metres wide and 6 metres high; it contains 35 Kt of copper ore, mainly malachite $Cu_2(OH)_2CO_3$ and azurite $Cu_3(OH)_2CO_3$. Before stacking onto the heap, the ore is pelletised into pellets of typical diameter 10–30 mm. The copper ore occurs in veins and grains in the pellets, and is accompanied by fines, binding agents, clays and silicates. The heap is continuously irrigated by a sulphuric acid solution. The acid percolates through the heap, and is collected by drainage pipes running beneath a gravel bed at the base of the heap. This acid solution, containing copper ions, then passes through solvent extraction and electrowinning processes to produce pure copper metal. Our emphasis is on the heap and the processes therein.

[1]CSIRO Mathematics and Information Sciences, Locked Bag 17, North Ryde NSW, Australia 2113. Email Noel.Barton@cmis.csiro.au.

[2]Mathematics Department, Tamaki Campus, The University of Auckland, Private Bag 92-019, Auckland, New Zealand. Email g.wake@auckland.ac.nz.

Further details of heap leaching have been given by Townson and Severs (1990), whilst West and Connor (1996) describe particular features of the work carried out by Nifty Copper Operations. Eriksson and Destouni (1997) describe copper leaching from mining wastes, with particular emphasis on chemical processes.

In order to understand the process, WMC has carried out laboratory and column tests. These gave copper recovery as a function of time and amount of acid application. The results are different to actual recovery results for the heap, indicating lack of understanding and inability to scale up from laboratory studies to the pilot study. A major thrust of this study is to explain this discrepancy, and in so doing, to give a predictive capability for the performance of the heap.

In Sections 2–6, we describe complementary approaches to address the above problem. In Section 2, we analyse the data and give empirical results for the performance of the heap. This gives insights into the importance of voids in the heap, and the need to manage their effect.

The chemical processes describing the action of acid on ore were thought to be relatively straightforward, and can be described by a lumped reaction. Basic equations are given in Section 3. In a full model, these equations would need to be coupled with equations describing fluid movement.

Section 4 examines individual pellets. Simple scaling arguments were applied to show how quickly copper is removed from ore veins by acid. This takes place on a timescale of days, and is an important factor in the performance of the heap. Over time, the pellets become moist and soft; this causes the heap to slump by about 15%. These pellet models provided the basis for empirical expressions for overall copper recovery.

In Section 5 and 6, we address the acid motions in the heap using continuum and probabilistic models. Gas (carbon dioxide) production is not considered in this model since we believe insufficient gas is produced to affect the dynamics of the heap. Overall, the process can be described as percolation. At pellet scales (i.e. 1–3 cm) a probabilistic model is used to simulate flow of acid around pellets thereby forming channels. 'Rivulets' or 'trickles' might be better descriptors. At the many-pellet scales (i.e. 1 m), we can describe the dynamics of the flow by continuum models embodying the percolation process. Spatial variation of the flow on this scale could also be described as channelling. The continuum model, described in Section 5, is crucially dependent on empirical expressions for permeability and potential functions for the flow. The probabilistic models, described in Section 6, give insights into pellet scale behaviour. They show channels on these scales form within 1 m of the top of the heap.

In the concluding section, we give conclusions and recommendations. We believe that more experimental and pilot tests need to be done. Further mathematical analysis, based on experimental results and a better description of the processes, is warranted. An important practical suggestion is to deploy instrumentation in the heap to measure flow properties and acid concentrations. More intensive irrigation strategies and physical intervention into the heap might improve copper recovery.

2. Data analysis and simple empirical models

2.1 WMC data and its analysis

WMC provided data for four tests as shown in Figures 1 and 2. These were for two laboratory scale tests, one test (NIF113) on a large column of ore, and the actual pilot study on the heap at Nifty Copper Operations. The challenge is to use data from the laboratory and column tests to predict the performance of the heap. A preliminary observation would be that it takes far longer (see Figure 1) and far more acid than expected (see Figure 2) to recover the copper from the heap.

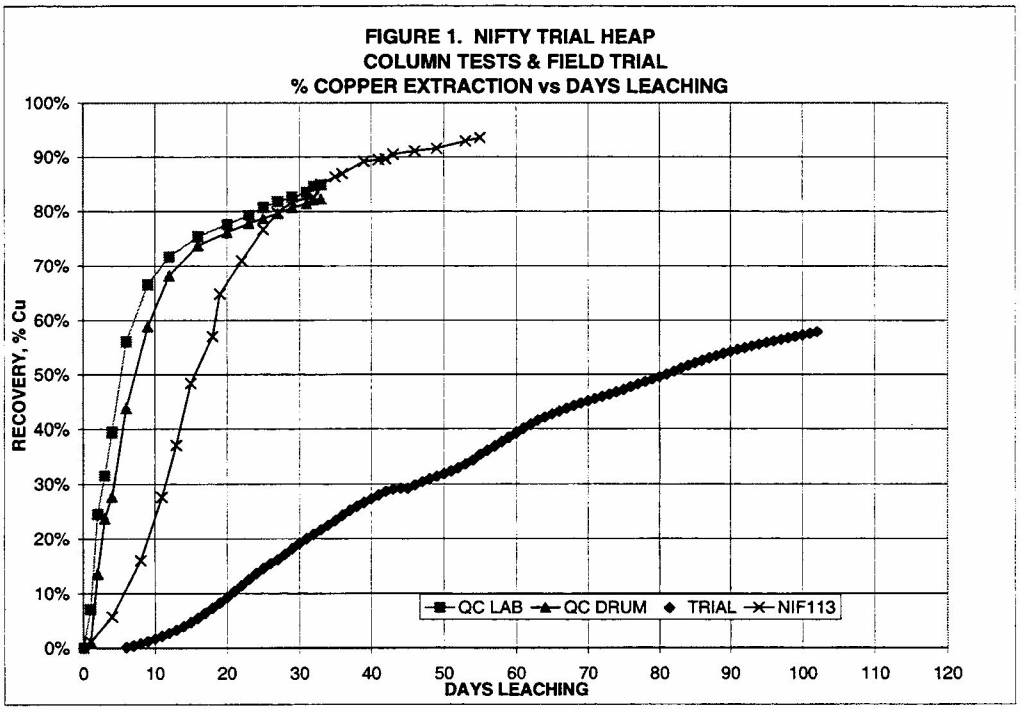

Figure 1: Cumulative Copper recovery (%) versus time (days).

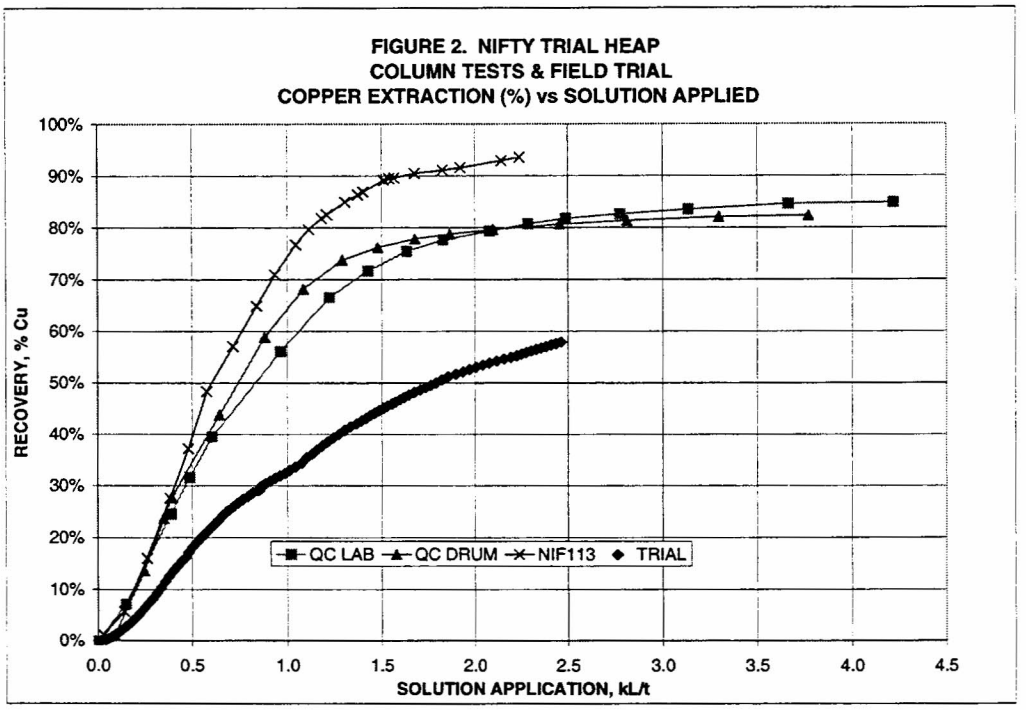

Figure 2: Cumulative copper recovery (%) versus solution application (kl/t).

Examination of the WMC data shows that the NIF113 data was obtained at an acid application rate which was twice that used for the pilot study. Given this information, one would not expect the same rates of copper recovery between NIF113 and the pilot heap. Neither would one expect that copper recovery should be dependent only on acid applied, particularly since we know from the previous section that the pilot heap is not saturated. We would however expect that NIF113 and the pilot heap should give similar results for copper recovered as a function of acid converted by chemical reaction. This is confirmed for the early stages of the data (Figure 3). At later stages, NIF113 shows that the amount of acid consumed exhibits irregular behaviour. This may be due to acid and/or copper ion retention in the column. Section 2.2 indicates a significant amount of liquor between the ore particles which makes accurate measurements difficult. Further the amount of liquor retained can be expected to change as the ore compacts. We note that the actual heap does not display this behaviour. Presumably this is because the heap is so large that many individual slumping events occur, and their effect is averaged out.

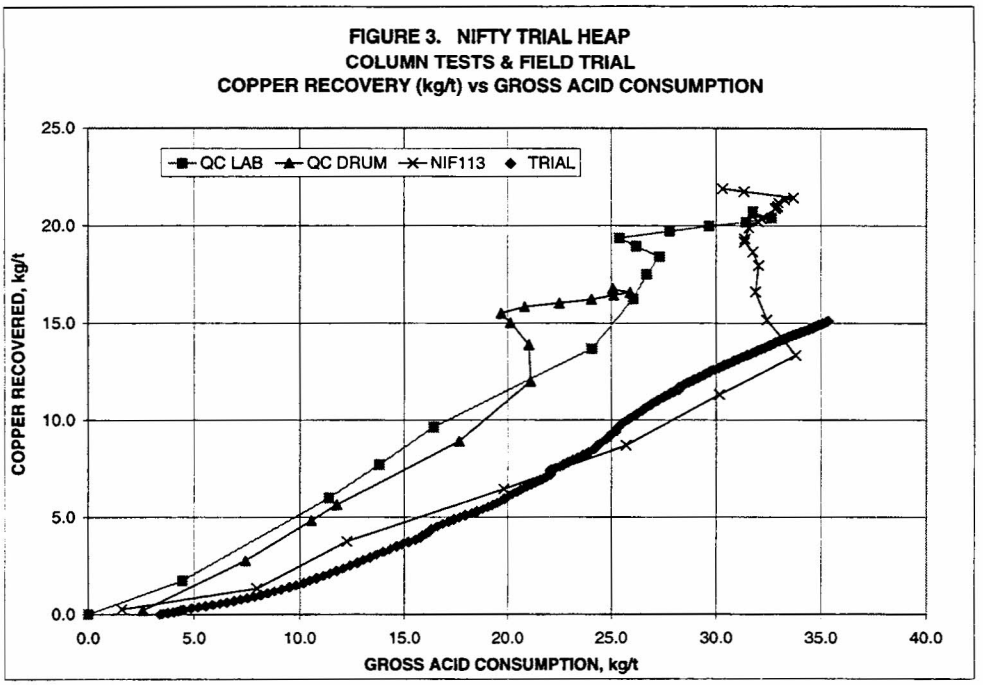

Figure 3: Cumulative copper recovery (kg/t) versus acid consumed (kg/t).

2.2 Amount of voids and acid

We are considering the flow of liquid sprinkled onto a heap of irregular pieces (pellets) of ore which average about two centimetres in diameter. A small drop of liquid contains about 0.05 ml, and the target flowrate of 15 l $hr^{-1}m^{-2}$ represents only one such drop every two minutes for each square centimetre of area, or about two drops per minute per pellet.

Assuming the pellets are spherical and uniform in size, the heap consists of about 60% pellets and 40% voids; similar proportions are found experimentally for irregular rocks. A mass balance on the cumulative inflow to the heap and the cumulative outflow shows an initial deficit in the outflow corresponding to about 20% of the voids being filled with acid solution in the steady state.

Hence the major portion of the heap is not saturated with acid although there might be a small saturated region at the bottom of the heap. We believe that the heap supports a slow flux of acid, which trickles from stone to stone, probably following well-defined paths marked out by the passage of previous drops.

2.3 Curve fitting approach

The aim of this subsection is to fit curves to the data displayed in Figures 1 and 2 so they could be used to predict copper output. The y-variable of these curves represents cumulative copper recovery (%) and is thus similar to a cumulative distribution function. Indeed, the density functions associated with these curves bear remarkable resemblance to log-normal curves. Thus, it was decided to fit a log-normal curve to the data. That is, we fit a curve of the form

$$y(x) = \int_0^x f(t)dt$$

where $f(t)$ is the probability density function of a log-normal distribution

$$f(t) = \frac{1}{t_\sigma\sqrt{2\pi}} \exp\left\{-\frac{1}{2}\left(\frac{\ln(t) - \mu}{\sigma}\right)^2\right\}.$$

In this expression, t is the acid solution application in kl/t. What is required is a two-parameter fit (for μ and σ).

Figure 4 shows the actual and fitted recovery rates for the four tests. The fit for the actual pilot heap is very good. The fitted values are $\mu = 0.675$ and $\sigma = 1.34$. As the operational heap will be similar in dimensions and structure to the pilot heap, the fitted curve can be used for predicting daily copper yield from an actual heap. At the time of the study, 75% of the copper has been recovered after 272 days. Using the fitted curve, one can predict that it will take 395 days to recover 85% and 890 days to recover 95%, the target of WMC, assuming the solution application is held at a constant rate equal to the average of the past.

This approach gives a result that is operationally useful, without requiring further experimentation to estimate various constants that are necessary in a mathematical modelling approach. It does not aim to provide any insight into the physical and chemical mechanisms that operate within the heap.

2.4 Two simple empirical models

Two sub-groups provided explanations for the behaviour of the cumulative yield curve in the heap. Their models, which are not mutually exclusive, are now described.

Model 1

The model represents the heap as a horizontal rectangular grid of m by n points. At time 0, k vertical channels are created in random locations through-

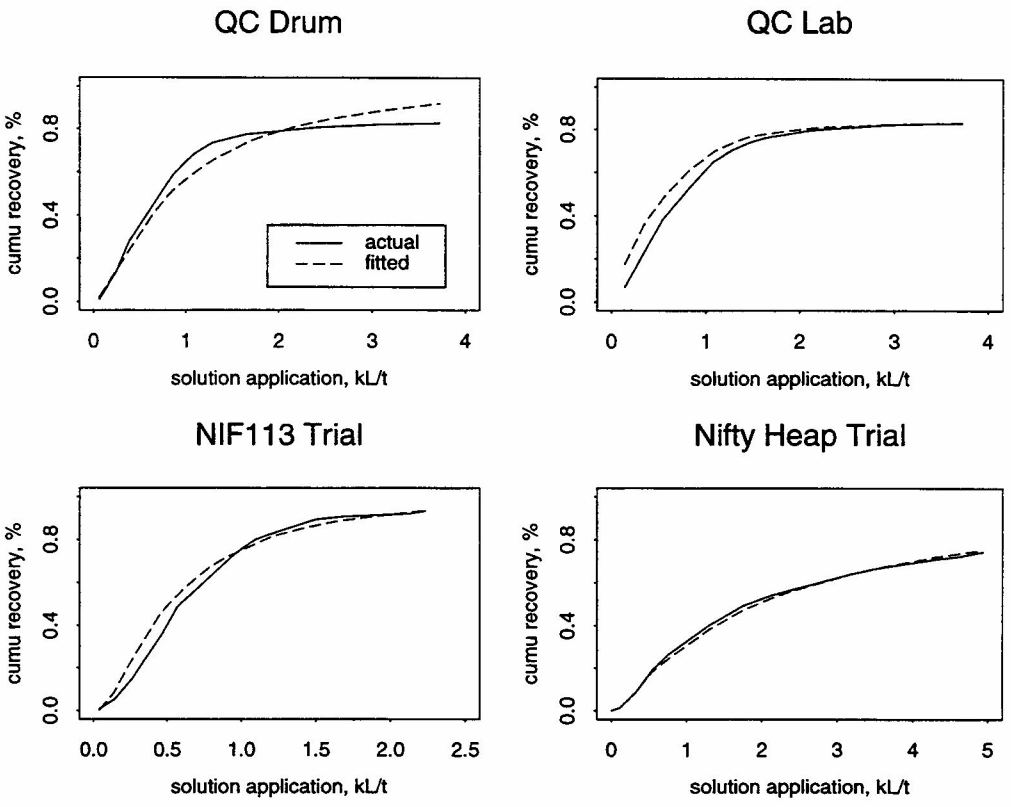

Figure 4: Fit of log-normal distribution to cumulative copper recovery from the heap.

out the grid. The flow rate of channel i at time 0 is $f_{i,0} = 1$. The starting concentration of copper at each point (x, y) in the grid is $C_{x,y,0} = 1$.

The following steps are then iterated in order:

Leaching. The output yield $Y_{t+\Delta t}$ of copper for time $t + \Delta t$ is computed, based on the locations and flow rates of the channels at time t. Computing the yield for $t + \Delta t$ rather than t ensures that the yield after zero time is zero (a non-zero yield after zero time is impossible). We put

$$Y_{t+\Delta t} = \sum_{i=1}^{k} Y_{i,t} + \Delta t$$

where $Y_{i,t+\Delta t}$ is the yield at $t + \Delta t$ due to channel i. The $Y_{i,t+\Delta t}$ in turn are found by assuming that each channel leaches a fixed proportion λ of the copper present.

That is

$$Y_{i,t+\triangle t} = \begin{cases} \lambda C_{x_{i,t},y_{i,t},t} & \text{if } f_{i,t} = 1 \\ 0 & \text{otherwise.} \end{cases}$$

The concentration at time $t + \triangle t$ is what remains after the leaching at time t, that is

$$C_{x_{i,t},y_{i,t},t+\triangle t} = C_{x_{i,t},y_{i,t},t} - Y_{i,t+\triangle t}.$$

Since the time that a channel stays in the same place will affect the proportion λ of copper that it leaches in that time, λ was made a function of $\triangle t$. The function used was derived directly from the leaching curves for the Lab and Drum tests provided by Nifty Copper, under the assumption that these leaching curves also apply to each channel within the heap. The values of λ that were tried were 0.1, 0.3, 0.5, 0.7 and 0.8, corresponding respectively to values of $\triangle t$ of 1, 3, 5, 11 and 24 days.

Channel movement. The channels are allowed to move randomly within the heap. That is, if channel i is in the interior of the heap at time t, its position $(x_{i,t+\triangle t}, y_{i,t+\triangle t})$ at time $t + \triangle t$ is computed by randomly adding -1, 0 or 1 to $x_{i,t}$ and $y_{i,t}$ with probability 1/3 for each possibility. However, if channel i is at the $x = 0$ boundary of the grid for time t, then $x_{i,t+\triangle t} = 0$ or 1 with probability 2/3 and 1/3 respectively. Other boundaries were treated similarly.

Merging of channels. If, at any time, two channels occupy the same grid point, they are merged. That is, if channels i and j with $i < j$ are in the same place at time t, $f_{j,t}$ is set to zero.

A reasonable fit to data was obtained using $m = 10$, $n = 10$, $k = 18$ and $\triangle t = 5$ days. The proportion leached as a function of time is shown in Figure 5. Reducing or increasing the number of channels had the effects, respectively, of reducing or increasing the rate of leaching until about 50 days had passed and 60–80% of the copper had been leached. The model does not however replicate the delay before the onset of leaching in the heap (about 5 days).

The results appear to show that leaching is a first order process with a rate constant that changes over time. The model shows that this might be explained by a reduction in the number of channels with time. If the channels in the model are allowed to pass through each other when they meet (that is, the number of channels is constant), the rate constant is invariant with time, and the yield curve is exponential.

Model 2

In this model, which is an extension of the model in Section 2.3, the cumulative recovery from the pilot heap is written as an integral which expresses the

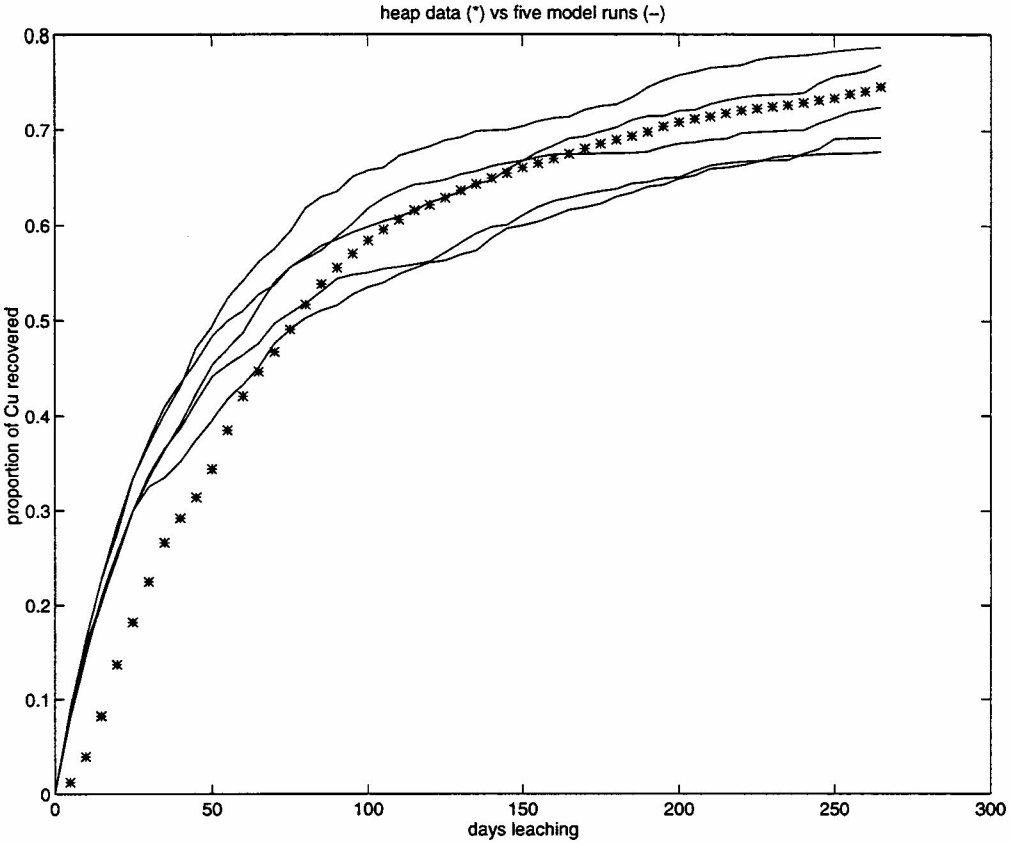

Figure 5: Simple empirical model for proportion of copper leached as a function of time.

total recovery from a number of columns with different total acid additions:

$$R = \int_0^\infty d(x, \mu, \sigma) r(x) dx$$

where $r(x) = 0.9(1 - \exp\{-1.4x\})$ is the proportion of copper recovered from a column (based on the laboratory experiments), x is the total amount of acid applied (in kl/t), and $d(x, \mu, \sigma)$ is the distribution of acid to the columns for a mean addition of μ and a standard deviation of σ. A log-normal distribution is a reasonable choice for the distribution of acid addition, that is

$$d(x, \mu, \sigma) = \frac{\exp\{-([\ln(x) - \alpha]/\beta)^2/2\}}{x\beta\sqrt{2\pi}}$$

with mean $\mu = \exp\{\alpha + \beta^2/2\}$ and standard deviation σ given by

$$\sigma^2 = \exp(2\alpha + \beta^2)(\exp(\beta^2) - 1).$$

The expression for R can be evaluated numerically for given values of μ and σ. Having chosen the mean μ, a value of σ can be determined to give values that correspond to the heap data.

The results for different mean addition rates can be summarised in the following table:

μ(kl/t)	0.5	1.0	1.5	2.0	2.5
σ	1.17	1.8	2.6	3.2	4.1
σ/μ	2.3	1.8	1.7	1.6	1.6

We see that the standard deviations of the flow distributions are larger than the means. This gives a distribution of acid additions that is well spread, with a major portion of the distribution of low addition rates, and a long tail going up into the high rates. A more uniform distribution of acid in the heap does not match the observed recovery of copper. It is concluded that the major proportion of the acid solution must flow in channels through the heap.

3. Chemical processes

The chemical processes involved, like many other situations, will consist of many chemical reaction stages. We expect that the chemical action of the acid on the ore can be lumped into one reaction which is unlikely to be rate-determining:

$$Cu_2(OH)_2\ CO_3 + 4H^+ \rightarrow 2Cu^{++} + CO_2 + 3H_2\ O$$

malachite + hydrogen ions $\rightarrow$ copper ions + gas + water

A similar outcome is likely with the action of the acid on azurite.

The rate-determining step will be the availability to the malachite of H^+ ions in the heap, which in turn is determined by the physical processes of percolation through the heap. A problem of considerable importance which is investigated in Section 4 is the process of the acid eating into the pellets. The ore itself is embedded in veins within the rock size pellets. This process could be rate-determining, but nonetheless the chemical processes could still be considered as the lumped reaction above, where the concentration of acid determines the rate of production of copper ions and ultimately copper metal. Accordingly, the chemical process can be modelled by a rate equation of the form

$$\frac{\partial}{\partial t}[Cu^{++}] = \text{const.}[H_2SO_4].$$

In a full model, this equation needs to be coupled with the fluid motion equations formulated in Section 5. In that section, the concentration of acid $[H_2\ SO_4]$ will

be closely related to $\theta(x,t)$ which is the moisture content expressed as volume fraction of water to the total volume. If we neglect the acid depletion, which is probably realistic in view of the fact that the key factor in the lower performance of the heap is the occurrence of dry areas, then the concentration of the acid will be proportional to $\theta(x,t)$.

With these simplifying assumptions the determination of θ from the fluid dynamics of the heap described in Sections 5 and 6 would directly give the rate of copper production in the above equation. Subsequently it may be considered pertinent to introduce more complex chemistry and this will mean more detail will be necessary. However, experience suggests the physics is the more important step, and our perception is that moisture, carrying the acid concentration with it, controls the rate of process. A recent paper by Eriksson and Destouni gives substantial details about chemical processes occurring in leaching of copper from mining wastes.

4. Pellet models

The Group took the view that processes at the individual pellet level were rate-controlling in terms of the performance of the pilot heap. We suppose that acid penetrates the pellets thereby removing copper. On dimensional considerations, the time T for acid to leach copper a depth d into the pellets is given by

$$T = d^2/\kappa$$

where κ is the combined diffusion coefficient of copper and acid within fissures in the pellets.

The cumulative amount of copper extracted from the pellet of diameter a is

$$\begin{aligned} p(t) &= \rho\{V_0 a^3 - V_0(a-2d)^3\} \\ &= V_0\rho\{a^3 - (a-2\sqrt{\kappa t})^3\}, \quad \sqrt{t} \le a/2\sqrt{\kappa}. \end{aligned}$$

Here ρ is the density of copper in the ore and V_0 is an ore shape constant. Now suppose the pellet size is distributed between $a = 0$ and $a = D$ (diameter) with mass fraction $m(a)da$ between a and $a + da$. The cumulative amount of copper extracted from particles in the size range a to $a + da$ at time t is

$$q(t) = \begin{cases} p(t)m(a)da, & 0 \le t \le a^2/4\kappa \\ V_0\rho a^3 m(a)da, & a^2/4\kappa \le t \le D^2/4\kappa. \end{cases}$$

The total copper removed at time t is

$$V_0\rho \int_0^D a^3 m(a)da - V_0\rho \int_{2\sqrt{\kappa t}}^D (a - 2\sqrt{\kappa t})^3 m(a)da$$

and the fraction of copper removed at time t is

$$C(t) = 1 - \int_{2\sqrt{\kappa t}}^{D} (a - 2\sqrt{\kappa t})^3 m(a) da \Big/ \int_0^D a^3 m(a) da \, .$$

This expression can be evaluated for different distributions $m(a)$:

uniform distribution: $m(a) = 1/D, \ 0 \le a \le D$

$$C(t) = 1 - \left(1 - \frac{2\sqrt{\kappa t}}{D}\right)^4$$

triangular distribution: $m(a) = 2a/D^2, \ 0 \le a \le D$

$$C(t) = 1 - \left(1 - \frac{2\sqrt{\kappa t}}{D}\right)^4 \left(1 + \frac{\sqrt{\kappa t}}{2D}\right)$$

As an illustration of the use of these results, suppose that the pellet distribution is uniform and that it takes one day for the copper to start leaching. Assume also that the eventual fraction of copper that is recovered is a constant, for example, 0.9. Then least squares fitting of the expression for $C(t)$ to laboratory data gives

$$C(t) = 0.9(1 - (1 - A\sqrt{t-1})^4)$$

where A = 0.0849. Figure 6 shows the results.

Lastly, the above arguments could be modified for the actual heap by using a distribution of columns. The implementation of this could be based on either of the models given in Section 2.4.

5. Continuum models of the flow

The Group considered various continuum models for the flow of acid through the pilot heap at Nifty Copper Operations.

One plausible model for flow of liquid through an unsaturated porous medium was discussed at the 1995 MISG (Fulford and McElwain, 1995). The dominant equation is mass conservation:

$$\frac{\partial \theta}{\partial t} = -\nabla \cdot \mathbf{J}$$

where $\theta(\mathbf{x}, t)$ is the moisture content expressed as volume fraction of water to total volume of the medium, and $\mathbf{J}$ is the volume flux of liquid. The volume flux $\mathbf{J}$ can be expressed using Darcy's Law as

$$\mathbf{J} = -K(\theta)\nabla(\Psi - z)$$

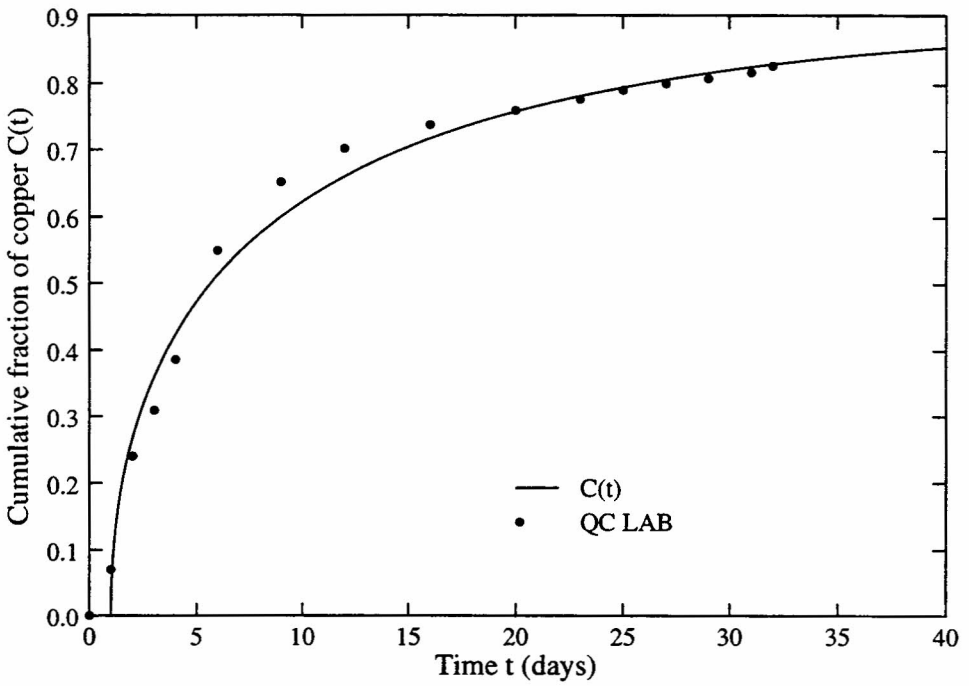

Figure 6: Leaching from a uniform distribution of pellets using dimensional arguments.

where $K(\theta)$ is the hydraulic conductivity, Ψ is the suction potential and z is the vertical distance from the origin, measured in the direction of drainage. The term $\Psi - z$ is the total pressure head, with Ψ due to capillary forces within the pores resulting in a suction force, and the remainder is the gravitational potential.

To solve the above equation for θ, it is necessary to provide empirical forms for $K(\theta)$ and $\Psi(\theta)$. There is a rich literature on this subject. For example, there is a commercial code LEACHM (available from C Vision Pty Ltd in Sydney) that simulates water and solute transport through unsaturated soil. Documentation on LEACHM provides models for K and Ψ in the soils context. For drainage from iron ore heaps, the following expressions were suggested in the 1995 MISG report:

$$K(\theta) = K_s \left(\frac{\theta - \theta_m}{\theta_s - \theta_m} \right)^3$$

$$\Psi(\theta) = \frac{1}{\alpha} \log \frac{K(\theta)}{K_s}.$$

Here K_s is the hydraulic conductivity for saturated flow, θ_s is the saturation

moisture content, θ_m is the water content that cannot be removed by drainage, and α is the sorptive number related to characteristic pore radius R by

$$\alpha = \frac{R}{7.4 \times 10^{-6}}.$$

For an iron-ore stockpile with $R = 1$ cm, then $\alpha = 1.4 \times 10^3$ m^{-1}. For fine soils, α is much smaller.

Some preliminary computations for the above equations were carried out during the MISG using the PDE solver *Fastflo*. The results were unexceptional and did not indicate channelling.

There are two particular difficulties with use of the above model as a computational framework for acid flow through the copper ore heap:

- The model assumes that the main driving forces of the flow are capillary forces within small pores and gravity. However, for the copper ore heap with relatively large pellets (2 cm), the dominant effects in the bulk of the heap are gravity and drop-to-drop trickling.
- Close to the bottom of the heap, the flow might be saturated, and this is not considered in the above equations.

To derive a better continuum model for the flow, it would be best to start with a conservation equation for moisture content

$$\frac{\partial \theta}{\partial t} + \nabla \cdot (\theta \mathbf{u}) = Q$$

and introduce plausible expressions for $\mathbf{u}$ and Q. In the unsaturated regions, $\mathbf{u}$ would be dominated by a mean component in the direction of gravity and stochastic component reflecting the drop-to-drop trickling. In saturated regions, $\mathbf{u}$ would be given by a conventional Darcy Law. It would be useful, although non-trivial, to derive equations for the mean of θ and the fluctuations from the mean. Such equations could only be derived from good experimental data.

Once such equations had been derived, we anticipate they could be readily solved using a PDE solver such as *Fastflo*. This would enable simulation of the drainage of the whole heap, and coupling to the chemical models as outlined in Sections 3 and 4.

6. Probabilistic models of the flow

Several probabilistic models were developed during the MISG. To set the scene for these models, recall the arguments in Section 2.1 that only two small

drops impinge each minute on average pellets near the top of the heap. These drops will have a strong tendency to trickle from pellet to pellet along constant paths. If every drop follows its predecessors exactly, it is easy to see that the drops will progressively combine into fewer and fewer channels. It was considered that mechanisms that could split a drop, such as landing on a sharp edge, would occur infrequently and hence they have not been considered in this analysis.

A mechanism like this can be represented as a matrix of layered cells dimensioned, say, 1 cm by 1 cm horizontally. It is reasonable to assume that a high proportion of the acid falling on each cell will flow to a cell in the next layer which is straight down or offset by one cell.

We estimate from simple geometrical considerations that a particle will on average deflect a drop sideways by about 0.3 times the particle's diameter. Treating these as random deflections, the drops move sideways by an average of $0.3\sqrt{N}H$ where N is the number of layers and H is the height of a single layer. For cells with a horizontal side of 1 cm, we therefore expect that a drop will move laterally by 1 cm for a vertical drop of 10 cm.

A 3D simulation was therefore carried out in which the cells have dimensions 1 cm $\times$ 1 cm $\times$ 10 cm. The bulk of the flow from each cell moves sideways by one cell in flowing down to the next layer; the remainder is spread evenly between the five cells below. No flow leaves the matrix at the sides.

Simulations were done for a matrix with 10 cells (0.1 m) on each side and 30 cells (3 m) deep. The figures in Figure 7 show the standard deviation, the skew and one tenth of the number of cells with a flow of less than 0.05. The average flow on each level is 1.0 in all cases. The figure of 0.05 represents a flow of about one drop per pellet every 10 minutes.

One simulation was run with 100 cells $\times$ 100 cells (1 m $\times$ 1 m) and a spread of 5%. There is evidence of slow growth in the results but overall the values do not appear to be significantly different from those in the earlier simulations.

To extract copper from a particular pellet, there must be an adequate flow of acid over the surface of that pellet. The simulations show that quite a high proportion of the stones will be relatively dry at any given time.

A second probabilistic model developed at the MISG was based on a variant of the probabilistic process by which acid flows from layer to layer. This second model, not presented here, also showed that channelling occurred after about 10 layers and was thereafter reasonably stable.

These probabilistic models raise several observations and questions.

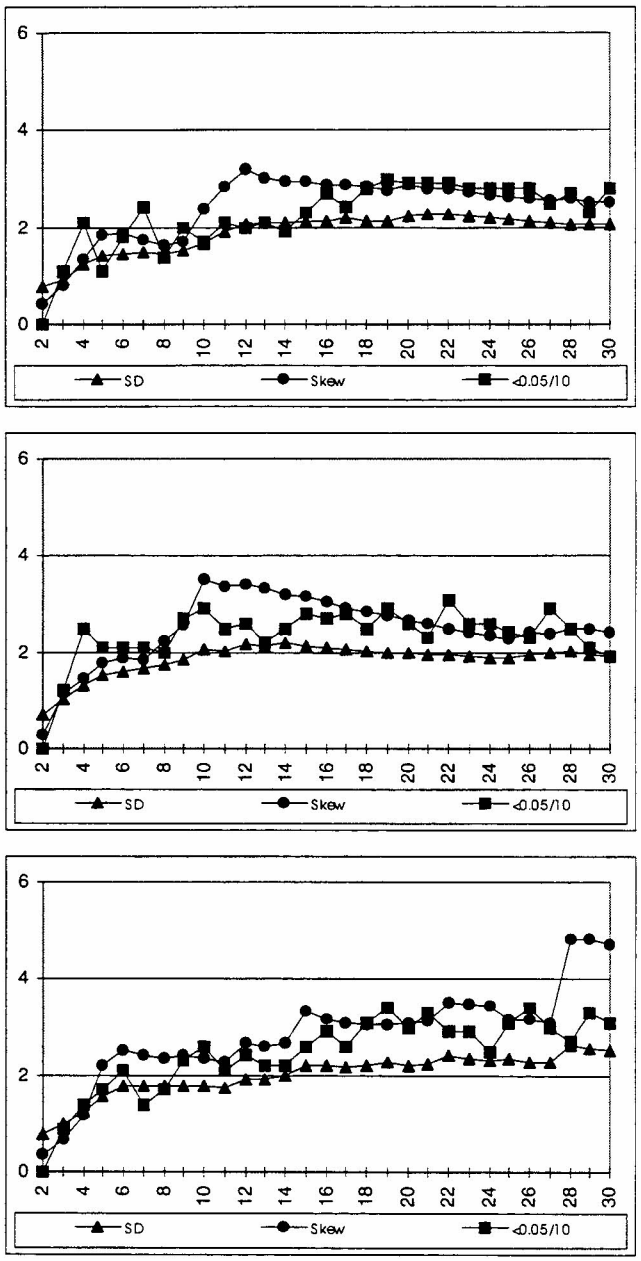

Figure 7: Three simulations with 5% of the flow being spread. Despite the differences between these simulations, the stable SD and skew, of around 2 and 3 respectively, are very similar, and reasonable stability is reached within about 10 layers (1m). About 30% of the cells have a flow of less than 0.05.

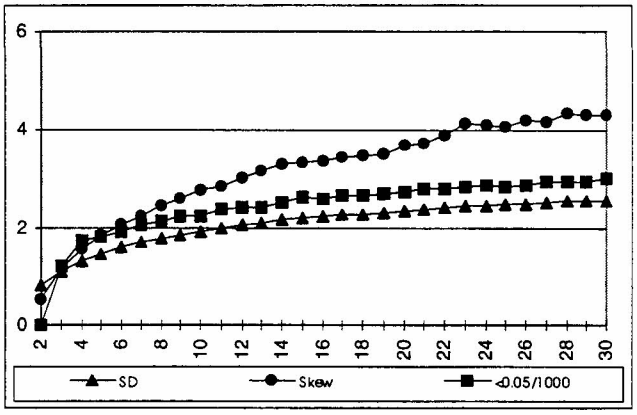

Figure 8: Simulation with 5% of the flow being spread for a 100 × 100 × 10 matrix.

- What is the physical scale of the channelling? Indications are that the channels are perhaps 10 to 30 cm apart, and they develop about 1 metre into the heap.
- How stable are the channels? Initial simulations suggest they are stable with perturbations to the probabilities changing only the few layers immediately below.
- What is the stationary solution? The process is Markovian and should have a stationary solution.
- The simulations give a plausible picture of what the flow looks like on the horizontal scale of 1 m × 1 m and a vertical scale of perhaps 3–5 metres. The current methods (based on the Excel spreadsheet) are computationally expensive, but huge speed gains could be expected if the simulations are programmed in C or FORTRAN. This should allow simulations up to, say 10 m × 10 m, which is significantly larger than the predicted horizontal motion of the channels and may be sufficient for practical applications.

7. Conclusions and recommendations

The Group considered several aspects of leaching of acid through the pilot heap at Nifty Copper Operations. These aspects include:

- simple data-fitting exercises to predict the overall performance of the heap;
- development of lumped chemical models within a continuum framework to describe the chemistry in the heap;

- development of a model to predict leaching from individual pellets and from distributions of pellets;
- consideration of continuum models for acid flow through the heap; and
- development of probabilistic models that described acid flow around collections of pellets on a scale smaller than that for the heap itself.

These various investigations shed light on the performance of the heap. The probabilistic models, in particular, helped to gain an understanding of the acid flow through the heap. These however give a simulation based on small scale events, and might not be directly applicable to the large scale simulation of heap performance.

The Group makes the following suggestions for Nifty Copper Operations to consider:

1. Carry out more laboratory tests for leaching, particularly for wider columns (or even over plates) and at higher flow rates.
2. Disturb the heap, for example by raking or shaking, to break up channels that form within a metre of the top surface.
3. Irrigate the heap in a less regular fashion, particularly through saturation irrigation in various parts, perhaps also by injection of acid into the bulk of the heap.
4. Instrumentation of the heap (and laboratory tests) to give more information on channelling and acid concentration.

To simulate the overall performance of the heap, the chemistry (Section 3), the leaching rate of particles (Section 4), and the flow of acid through the heap (Sections 2.4, 5 and 6) need to be combined. It is currently not clear what type of flow model should be used and possibly a combination of techniques will give the best results. We see the development of this model as a large activity which would require substantial experimental backup. We do however expect that the development of such a model would provide substantial benefits to mineral processing companies operating leaching heaps.

Acknowledgements

Work on this topic was managed by Noel Barton and Graeme Wake. A large and active team contributed to the discussions. It is appropriate to highlight particular contributions by Bill Whiten, Meei Pyng Ng and Tim Svenson

(Section 2), Charles Johnson (Section 3), John Hewitt (Section 4), and John Henstridge and John Whittle (Section 6). The WMC representative, Jeff Connor, enthusiastically kept the Group on track and provided valuable data and experience.

References

N. Eriksson and G. Destouni, "Combined effects of dissolution kinetics, secondary mineral precipitation, and preferential flow on copper leaching from mining waste rock", *Water Resources Engineering* **33** (1997), 471–483.

G. Fulford and S. McElwain, "Moisture movement in bulk stockpiles", *Proceedings of the Mathematics-in-Industry Study Group* (John Hewitt, editor, 1995), 48–65.

B. Townson and K.J. Severs, "The solvent extraction of copper — a perspective", *Mining Magazine* (1990).

J.N. West and J.W. Connor, "Nifty: Hydrometallurgy in the sand dunes", *AusIMM Annual Conference* (1996), 181–183.