Mathematics in
Industry
Study
Group
2009

Mathematics and Statistics in Industry Study Group 2009

Industry Project Reports: Equation Free Summaries

University of Wollongong

ANZIAM

Photo on front cover, from left: Prof Tim Marchant, MISG 2009 Director, Dr Maureen Edwards, MISG 2009 Assoc Director, Dr Maria Cozijnson, Bluescope Steel Research and Prof Alistair Fitt, Pro Vice Chancellor University of Southampton and invited speaker.

INTRODUCTION

Some of the project team for the ICT Research institute project

The Mathematics and Statistics in Industry Study Group (MISG2009), was held during the week 27th-31st January 2009 at the University of Wollongong. MISG2009 was hosted by the School of Mathematics and Applied Statistics at the University of Wollongong and ANZIAM, the Australian and New Zealand Society for Applied and Industrial Mathematics. The event was directed by Prof Tim Marchant and Dr Maureen Edwards. Administrative support was ably supplied by Ms Joell Hall.

Five industry projects were presented from Bluescope Steel Research, Geoscience Australia, Integral Energy and the ITC Research Institute. The projects spanned the areas of Applied and Financial Mathematics, Statistics and Operations Research. Bluescope Steel, Geoscience Australia and Integral Energy have all participated in previous MISG's and we are very thankful for their long-term support of the MISG concept.

100 people attended MISG2009 including 20 postgraduate students. As the AMSI Summer School for Honours and Coursework Masters students was held at UOW concurrently with MISG2009, a novel feature of the meeting was the Summer School student participation in MISG2009 as part of their studies.

MISG2009 was opened by Prof Joe Chicharo, DVC-International, University of Wollongong and closed by Dr Bob Anderssen, CSIRO. The invited speaker was Prof Alistair Fitt, PVC-International from Southampton University, UK. Dr Chris Breward, University Of Oxford and Dr Bob Anderssen, CSIRO also gave talks in our seminar session. Our thanks go to all of these people and, of course, to the project moderators. The moderators take responsibility for the industry projects and put in an inordinate amount of time and effort. These contributions are critical to the success of MISG.

The Australian Mathematics Sciences Institute (AMSI) supported postgraduate attendance at the event. Thanks go to AMSI, and to the University of Wollongong, for their financial support of MISG2009.

MISG2010 will be held at RMIT University, Melbourne 7-12th Feb. 2010 with A/Prof John Shepherd as Director. We hope that everyone that attended MISG2009 will be able to make it to Melbourne for MISG2010 and we wish John well with his plans for MISG2010.

In the following pages non-technical outlines of the projects, and the progress made, are presented. The full proceedings of MISG2009 will be published in late 2009

Tim Marchant,
Maureen Edwards,
Directors, MISG2009.

Project discussions for Bluescope Steel project on coating deformations.

Index

Slumping of steel coils

Bluescope Steel Research

Industry Representative:
Dr Maria Cozijnsen, Bluescope Steel

Problem Moderators:
Dr Bob Anderssen, CSIRO
A/Prof Mark McGuinness, Victoria University of Wellington

Student Moderator:
Ms Roslyn Hickson, UNSW@ADFA

Bluescope Steel manufactures, stores and transports coils of steel strip. They asked the Study Group to consider the slumping of these coils, where the coil cannot hold up its own mass and maintain the integrity of its cylindrical shape when it is stacked on its side. In particular, the inner hole of a slumped coil is distorted from a circular cross-section, and subsequent handling of the coil is impeded.

The steel strip has typical dimensions of thickness 0.3mm, width 500mm and length 1-2 kilometres, and is stored in a coil with typical dimensions of inner radius 250mm, and outer radius 750mm, laid on its curved side.

If the coils fail to maintain their circular cross-section, they become unusable and must be scrapped at considerable cost. This happens rarely, but even minor slumping is a major concern. The mandrels, onto which the coils are loaded at various stages of their processing, cannot be inserted if the inner hole is too distorted.

The major challenge posed to the study group was the development of a comprehensive understanding of the initiation and subsequent evolution of the slumping. In particular, Bluescope Steel was interested in quantifying the effect of coil mass and sheet thickness. They were also interested in the ability to predict the final shape of the coil cross-section, which from their experience tends to be either elliptical or triangular.

Though there is a comprehensive literature about coil winding, rewinding, optimal tension and more, it was of only marginal use for the current investigation. The specific type of problem that Bluescope Steel wanted investigated had not been previously analysed.

The situation confronting the MISG participants was an excellent example of one where the problem description was simple, clear and unambiguous, but actually getting a hook into the problem was quite challenging. A number of approaches to the problem were made – balance the forces in a localized representative increment; determine and then minimize the total energy contained in a wound coil (gravitational, bending, stretching and frictional energies) to find the related Euler-Lagrange equations; for a simpler system determine the total energy and corresponding Euler-Lagrange equations in order to obtain a better understanding of the underlying mathematics and insight about the nature of the problem; concentrate on an appropriate linearization and solve that first; treat it as a classical continuum mechanics problem and attack it directly; understand the role of the frictional forces on inter-layer slip in maintaining the shape of a wound coil; and analyse the interplay between gravitational and bending energies of a simple system to understand how a coil supports its own mass.

How tightly a coil is wound is critical. Too tightly, and a different kind of slumping occurs, a kinking due to excessive stresses at the inner wraps of the coil. Too loose, and each wrap of a coil acts independently and slumps gently since it cannot hold up its own weight. In between, and friction links adjacent wraps to strengthen the structure, effectively thickening wraps so that they can hold up their own weight.

At various stages of the deliberations all of these options were vigorously discussed. It was a necessary part of the brainstorming process required to generate useful insight. Some of the above options were investigated carefully and methodologically. On the one hand, explicit formulas for the gravitational and bending energies for a variety of simple configurations were obtained and utilized. On the other hand, insight was obtained about the tensorial nature of a stretching energy calculation.

These deliberations were of immediate value to the Bluescope Steel representatives as, by being directly involved in the discussions, it gave them the opportunity to enhance their own understanding about the nature and challenges of the interesting problem that coil slumping represents.

Coating deformation in the jet stripping process

Bluescope Steel Research

Industry Representatives:
Dr Cat Tu, Bluescope Steel
Dr Daniel Yuen, Bluescope Steel

Problem Moderators:
A/Prof Graeme C. Hocking, Murdoch University
Dr Winston L. Sweatman, Massey University

Student Moderator:
Ms Melanie E. Roberts, University of Western Australia

During the production of steel strips, the steel surface is usually coated with a layer of metallic alloy (e.g. zinc/aluminium) to protect it from corrosion. A mechanism for achieving this is the “continuous hot-dipped galvanizing process" for which the steel strip is passed through a bath of the molten metal coating. The steel surface is protected from corrosion prior to entry into the bath. After passing through and being coated by the molten metal the strip is pulled upward until eventually the metal coating cools and solidifies.

The thickness of the metal coating is controlled by a pair of air knives on either side of the ascending steel strip. Each of these air knives fires air at a high pressure along a horizontal line across the rising strip. The air forces some of the molten metal to return downwards into the bath. Above the air knives there is no further loss of metal coating material.

During recent developments in the production process, problems with the surface quality of the coating have arisen, in particular with some new advanced alloy coatings. Associated with each set of process conditions, such as the speed of the strip and the alloy used, there is an air knife pressure below which the coating is satisfactory but above which coating defects may appear. The defects take the form of waves, lines and pocks. The pocks are the most serious of the deformations as they are the deepest and correspond to a substantial thinning of the coating and dramatic reduction of the corrosion protection. Each pock has dimensions of the order of millimetres, is broader in the horizontal direction, has a downstream bump and has a depth of 70-90% of the usual alloy coating thickness.

The MISG group modelled the process as a two-dimensional, uni-directional fluid flow and derived an equation for the dynamics of the coating layer including all processes that they thought might be of importance. The magnitudes of the different terms including surface tension, shear, heat transfer, air pressure and gravity were estimated from typical parameter values to determine their relative importance. In the absence of the air knife, gravity was found to play no significant role. Surface tension was found to have only a small effect while the air pressure and shear appeared to be the dominant terms. The resulting “simplified" equation is a first-order partial differential equation, the solution of which turns out to be rather delicate depending on the upstream flux condition. This equation was consistent with earlier results described in the literature and the model used by Bluescope Steel.

The group also considered the effect of the formation of an oxide film on the surface of the molten coating. This seems to affect the surface properties but not the macroscopic fluid flow. The thermal conduction problem was also considered: heat transfer did not appear to be important over the length scale of interest.

Possible mechanisms for pock formation were discussed. Dissolved bubbles in the molten coating seemed unlikely, but the possibility of entrapment might be investigated by experiment. Contamination of the metal in the bath or in the air is already being investigated by the company. Vapour explosions were ruled out by the industry representatives. However, it seemed plausible that impacts of particles, entrained in the air knife, might produce a fracture in the oxide film. This is to be further investigated by calculation.

Future work will be used to confirm and refine some of the preliminary results. It would be desirable to find surface trajectories for more realistic pressure/shear terms. Using the full partial differential equation could show how an initial distortion will evolve: it could be used to consider whether small ripples could cause the line/wave defects. The techniques for finding the upstream flux condition can be further refined.

Provenance of sedimentary rocks

Geoscience Australia

Industry Representative:
Dr Keith Sircombe, Geoscience Australia

Problem Moderators:
Dr Geoff Pritchard, University of Auckland
A/Prof Ken Russell, University of Wollongong
Prof Bill Whiten, University of Queensland

Student Moderator:
Mr Paul Moloney, RMIT

Background

Geoscience Australia is an agency of the Australian Government. It provides “geoscience information and services such as maps, earth monitoring and strategic data about mineral and energy resources". It is very interested in the science of geochronology (analysing the age of geological events), which leads to a need to date rocks. This project is specifically interested in the dating of sedimentary rocks, such as sandstone, because these have been formed from the eroded detritus of other rocks and can thus provide a snapshot of the geological environment at the time when the detritus was deposited. The ages of uranium-bearing zircon grains in the sediment can be used as a proxy for the age of the rock from which the minerals originated, giving a guide to the provenance of the sedimentary rock. As zircon crystals form in molten magma they absorb uranium but not lead. Over time the uranium slowly decays to lead providing a means to determine the age of zircon grains. However, the zircon grains whose ages are determined are only a small sample from all the grains in the deposit, and although great care is taken the sample may have been contaminated by grains of other ages. Three challenges have arisen:

- How many analyses must be acquired for a sample to be representative of the population? Is true representation achievable?
- Can age distributions be objectively compared?
- Can the detrital ages limit the age of the sedimentary rock itself?

To aid investigation, Geoscience Australia provided a data set of 47 samples, comprising age determinations for a total of over 2500 grains of zircon.

The age of a single grain

The dating of individual grains, although not a formal objective at MISG2009, was briefly considered. It can be done by at least three different methods using the ratio of : $^{207}Pb/^{206}Pb$, $^{207}Pb/^{235}U$, and $^{206}Pb/^{238}U$. Each method produces an age with an associated uncertainty. A typical rule of thumb is to rely on $^{206}Pb/^{238}U$ for ages less than 1000Myr, and $^{207}Pb/^{206}Pb$ for greater ages. This could perhaps be improved upon by taking a weighted combination of all three estimates (having regard to correlation between them). Where different methods give widely disparate ages, the grain is labelled “discordant" and ignored in subsequent analysis.

The final age determination has an associated 1-sigma uncertainty which varies greatly between grains; it is anywhere from 1 to 200Myr in the data provided. The uncertainties tend to be greatest for ages around 1000Myr, and less for much younger or much older grains.

Estimating an age distribution

The ages of the grains in a sample can be used to estimate the age distribution of the grains in the geological source material. The sedimentary deposits typically derive from multiple sources of different ages, thus we want to estimate the whole shape of this age distribution, rather than only a summary statistic such as the mean. Two of the key questions are:

1. How many grains are needed to estimate the age distribution adequately?

2. How to compare or classify samples from different locations, on the basis of their age distributions?

Both questions imply the use of a metric to quantify the “distance" between two distributions. For the first, the term “adequately" implies that the estimated distribution is acceptably close to the true one. For the second, we need some objective measure of how similar one distribution is to another. There are several ways to do this. Regardless of the estimation method used, the number of grains needed for an adequate estimate of the age distribution will depend on the variability in that distribution and the type of distribution properties of interest. The more variable the data are, the larger will be the sample needed to be representative, and it is not possible to predict the required sample size until one has some knowledge of the variability. Perhaps the question could be re-formulated as “are there enough grains in this sample yet?".

Three approaches to defining the distribution of zircon particle ages were considered and are outlined in the following paragraphs. Two more paragraphs then describe uses for these distributions.

Empirical distributions. One method is to convert the age data directly into a discrete probability distribution with probability 1/n at each of the n ages in the sample (the “empirical distribution"). This is appealing because it requires no assumptions as to the nature of the true age distribution being estimated. The corresponding distance measure must be able to cope with both discrete and continuous distributions: the Kolmogorov-Smirnov and Wasserstein metrics are two ,rather different, possibilities.

Kernel density estimation. A second method is to estimate the shape of the density at a particular age by taking into account the number of estimated ages that lie within a close neighbourhood of that given age. This gives rise to what is known as “kernel density estimation". A closely related method is to use the estimated variability of each estimate to construct a smoothed density curve that is less jagged than we would obtain if we just used each value to produce an estimated density. With either method, we need only define a metric on probability density functions. This approach is pursued in [1], where the L_2 metric is used.

Parametric approach. A third approach regards the sample as having come from a mixture of several components with Normal (‘bell-shaped") age distributions; e.g., 25% of the grains have one particular age distribution, and the remaining 75% have another age distribution with a different mean. One needs to estimate the number of different Normal distributions that have been mixed, their means and variances, and the proportion of each component that is present. This approach has the advantage that the clusters found should divide the sample into groups of ages that correspond with the geological sources that make up the deposit.

The parametric approach allows estimating the probability of not finding a cluster consisting of a given proportion of the grains present. Also it is possible to test whether a small number of grains close together in age, is sufficient to claim a cluster exists in the data. For the classification problem, an appealing approach may be to pool some related samples to estimate the number and parameters of the normally-distributed components present, and then estimate the mixing proportions for each sample separately to determine the differences between the samples. These mixing proportions provide geologically useful information on how samples have varied with time or location.

A measure of dissimilarity between samples can be based on the mixing proportions, for instance the Euclidean distance or a test statistic similar to the Chi-squared goodness-of-fit test.

Relations between samples. Measures of dissimilarity from the analysis of the sample age distributions can be used as the input to standard clustering algorithms. These arrange the samples into a structure, known as a dendrogram, that shows which samples are related at different levels of closeness. Initial testing indicates that similar results are obtained from several different clustering algorithms. However different measures of dissimilarity give different clusters. It will be a matter of experience to determine which measures give the most useful geological information.

Age of sedimentary deposits. Given the ages of the zircon grains in the sample, the age of the sedimentary deposits is clearly later than the most recent age. Additional information on the likely time between the formation of the zircon grains and the formation of the sedimentary deposit, and on the range of possible ages for the sedimentary deposit, can be added as prior distributions. A small probability for the possibility of a younger outlier can be included. Using the grain ages and these priors, the distribution of ages for the sedimentary deposit can be calculated. This distribution will be very wide unless the priors provide sufficient information to narrow the range of ages.

Conclusions

The distribution of the ages of zircon grains found in a sedimentary deposit can be examined in several different ways. Analysis as a mixture distribution has the advantage of giving results that have a geological interpretation. Measures of the difference between distributions can be used to determine patterns of relations between samples. It is not clear which measures give the most meaningful geological results. The age of a sedimentary deposit, other than its upper limit, is not well determined by the zircon ages. Prior information can assist in the determination of age, but the result relies on the accuracy of the prior information.

References

[1] Sircombe K. N. and Hazelton M. L., Comparison of detrital age distributions by kernel functional estimation, Sedimentary Geology 171 (2004), 91-111.

VaR for Electricity Derivatives

Integral Energy

Industry Representative
Ms Nadima El-Hassan, Integral Energy

Problem Moderators
A/Prof Jeff Dewynne, University of Wollongong
Dr Otto Konstandatos, University of Technology Sydney

At regular half-hourly intervals, in the future, electricity retailers have to buy an unknown amount of electricity, the load, at an unknown price, the pool price, and sell this on at a known price, the transfer price. At each half-hourly interval, the unhedged earnings are the product of the load and the difference between the transfer and pool prices. This involves two random variables, the load and the transfer price, and therefore has a risk associated with it (in practice, this risk is quite large).

In order to reduce this risk, electricity retailers may use derivatives to hedge their earnings. Typical hedging instruments are caps and swaps, which come in a variety of flavours (flat, peak and off-peak) and tenors (monthly, quarterly and annual). These are traded on the Futures Exchange. More complicated instruments exist, but these tend to be relatively illiquid `over-the-counter' products which may be difficult to sell quickly should the need arise. We did not consider over-the-counter products.

In addition to their use in hedging a retailer's exposure to load and pool-price, derivative contracts may also be held in a so-called 'strategic portfolios'. Contracts in such portfolios are not used to directly hedge risk, although in some circumstances they may be transferred into the hedging portfolio.

The most important difference between the derivatives in the hedging portfolio and those in the strategic portfolio is that those used to hedge are held to maturity, and so their payoffs are realised, while those in the strategic portfolio are often sold before they generate any payoff. This implies the appropriate regulatory risk-measure for the strategic portfolio is *Value-at-Risk (VaR),* whereas for the hedging portfolio, it is *Earnings-at-Risk (EaR).* The basic problem in computing VaR is estimating the probability density for future values of a portfolio (conditional on the current state of the portfolio). This allows us to obtain the quantile of the portfolio's future value that corresponds to a particular level of

VaR. It is important to note that value here means *market value,* the value that could be realised if the portfolio were to be sold.

In general there are three generic approaches used to compute the VaR of a portfolio:

1. Analytic VaR, which usually amounts to assuming returns on all basic components of a portfolio are multivariate normal, leading to analytic formulae, which are notoriously inaccurate for portfolios with many derivative securities;
2. Monte Carlo (or parametric) VaR, in which one assumes certain distributions for the returns on the basic components of a portfolio and then runs large simulations to determine the probability density of future portfolio values; and
3. Historical (or non-parametric) VaR, in which one uses historic values for the returns on the basic assets in the portfolio in order to estimate likely future changes.

It is important to note that in most applications of VaR, changes in the prices of derivatives in the portfolio are computed from changes in the prices of the basic, underlying assets via a pricing model rather than directly. This assumes, of course, that we have a reasonably reliable algorithm for computing each derivative's market price as a function of its underlying assets' prices (for example, a Black-Scholes formula).

It is unlikely that this approach will work in our case. There are reasons for believing that no such formulae exist, or if they exist the input parameters required are not publicly available:[1]

1. The writers of electricity derivatives do not, and can not, delta-hedge because the underlying asset, electricity, can not be easily or efficiently stored;
2. The market for pricing electricity derivatives is not liquid;
3. *Theoretical* arbitrage opportunities are not necessarily *exploitable* because of the illiquidity and the inability to store the underlying commodity (electricity);
4. As a result of the small number of market participants, attempts to sell a large volume of derivatives over a short time may have an enormous (and adverse, from the seller's point of view) impact on the market price of those derivatives.

[1] For example, a hydro-electric generator may have to take into account its commitment to deliver contracted volumes of water to irrigators when deciding how many megawatts it can safely deliver when writing and pricing a cap contract.

Moreover, the relative illiquidity of the market means that even if a theoretical formula did exist there is no mechanism (arbitrage) to force market participants to use it, meaning that it is useless for the purposes of computing VaR (which is always concerned with *actual, realised portfolio values, under the real-world measure*).

This is not to say that it is impossible to find a model that gives a formula for the value of a derivative contract to a particular party, but there is no reason to suppose that anyone other than that party would necessarily agree with, or pay, that price. That is, there is no reason to suppose any such formula represents a *market* price.

This means that the traditional VaR approaches are not directly relevant to portfolios of electricity derivatives.

We decided that because the contractual features of the traded derivatives involved have not changed greatly over the last six years, and the pool-price's statistical features have not changed greatly over this period either, it is *reasonable* to use historical returns for derivatives prices, rather than attempt to find a pricing formula and then proceed in the traditional way. This also has the advantage that the rather large bid-ask spreads on the derivative contracts are automatically built into the estimate.

In view of the potential market impact of any attempt to liquidate a large portfolio, in order to estimate the distribution of returns on a derivative we need to know how we intend to sell it. We considered a variety of simple models for the way in which the illiquidity of the market affects the distribution of profit and loss if we attempt to sell off a portfolio of swaps or caps; the issue here is that if we wish to liquidate a portfolio over, say, a 10 day period then we can do it in many ways.

The most obvious is to sell it off all at once, say on Day 10. The problem is the market is so small that this may significantly decrease the prices of the derivatives we are trying to sell. It may not, but often it will.

An alternative is to sell off the portfolio in small pieces over the 10-day period, say 10% of the initial portfolio each day. Using historical data, this results in a quite different distribution of profit and loss to selling off everything at once (assuming that this has no effect on the market price).

A third alternative is to attempt to model the fact that some days you can sell large numbers of derivatives without moving the market, but often you can't. We modelled this with a Poisson process. For each day there is a (small) probability that selling off the whole portfolio will not move the market. If we try to sell and it doesn't move the market, we do so. If it does, we back off, sell a small fraction (say 10% of the initial portfolio if we have a 10-day horizon and try again the next day. We repeat this until the final day.

What we still need to do

Determine which (if any) of these three models best suits Integral Energy's needs (and, if appropriate, calibrate the third). Test the observed distribution of derivative returns for stationarity (or otherwise). Look at the joint distributions of the most common derivatives' returns and look for correlations (or analogous measures) between pool-prices, forward-prices and derivative-prices.

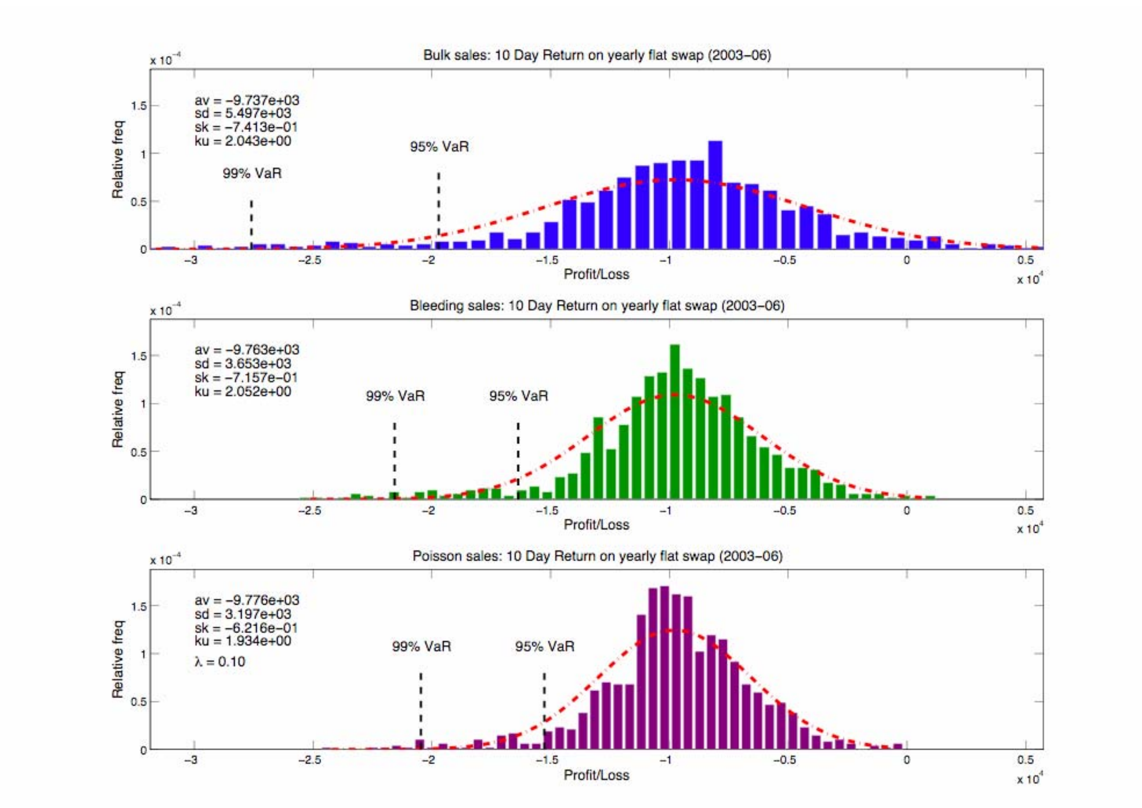

In this figure, the first frame shows the distribution of 10-day profit and loss assuming the entire portfolio of flat swaps can be sold on day 10 without moving the market price for the flat swap (the empirical profit and loss is computed using independent 10-day returns from 2003 to 2006). The second frame shows the distribution of 10-day profit and loss assuming that 10% of the initial portfolio is sold each day over the course of the 10-day interval (again, assuming that this has no impact on the market price each day and using the same historical data as the first frame). The third frame shows the distribution of profit and loss assuming there is a 10% chance that selling the entire portfolio has no market impact (in which case it is sold immediately) and a 90% chance that it does (in which case only 10% of the initial portfolio is sold and it is assumed that this has no market impact, the remainder is then treated in the same way on following days until the final day when it is assumed that it is possible to sell off whatever remains of the portfolio without impacting the market price). Note that the figure in the third frame represents an average over 10,000 simulations of the Poisson sale process and that the 10% probability of not moving market is unrelated to the choice of 10% of the initial portfolio as a fall back sale. The negative average profit is a result of the relatively large bid-offer spread on these contracts.

The red dashed lines indicate the normal density with the same mean and variance as the associated historical distribution and the 99% and 95% VaR levels (that is, the 1% and 5% left-quantiles) are computed from the historical distributions.

Effective communication in a virtual dungeon

ICT Research

Industry Representative:
Prof Farzad Safaei, ICT Research Institute

Problem Moderators:
Dr Philip Kilby, NICTA
A/Prof Desmond Lun, University of South Australia

Student Moderator:
Ms Giang Nguyen, University of South Australia

Back near the start of the internet, geeks played role-playing games with each other in what were called “multi-user dungeons”. They communicated with each other (by typing) and shared a common environment (the dungeon).

Those dungeons have now evolved into the virtual worlds of *World of Warcraft*tm and *Second Life*tm, where you are represented by your avatar in a completely digital world. You can wander around the virtual world seeing the sights, hearing the sounds, and interacting with other people through their avatars. This “interaction” can be as simple as typing a message, but can also include speaking and sharing audio, sharing pictures and video, and even live video feeds (yes, you can sit and watch a movie together). You might be chatting at an intimate party with a few other friends, or you may join couple of thousand others to watch a live concert in a virtual park.

All this interaction means lots of data that has to be carried around the computer network very quickly. And while two people chatting may be close in a virtual world, they may be physically separated by half a planet. Yet the sound bytes have to arrive within 100ms, or the delay will start to be noticeable, and detract from the experience.

Unlike the internet, where data is free to roam, these sorts of systems are designed so that most of the traffic is sent directly between the computers of people who are actually in the virtual world at the time. But because of limitations of those computers, they can only connect to a few others at anyone time – even though, if they are hosting a concert, they need to get information to thousands of others.

That is the nub of the problem examined during MISG: how do we connect servers in a virtual network so that the limits of each server are respected, but still ensure that people in the virtual world can communicate effectively.
One thing that makes it a bit easier is that "effective" communication gets easier with (virtual) distance. While I need to communicate with someone standing in front of me very quickly, I don't need the signal from the next room for a while. And if something happens a long way away, that information may never have to reach me.

What makes it harder is that things change quickly in the virtual world. I can decide I don't like Sweden and that I'll check out Borneo instead. So the people I am close to (virtually) can suddenly change quite a lot. This requirement, plus the fact that we may have to design a communications network for many thousands of interacting users, means the solution techniques must be very quick.

During MISG we looked at two methods: a quick method, called a "heuristic" and a slow exact method. The quick method lets us solve the problem in the time required, but cannot guarantee that the solution is very good. The exact algorithm ensures we find the very best solution, but cannot work quickly, or solve the problems with thousands of users in the real world. Mathematicians like these exact methods, however, because it lets us see how good the heuristic methods really are. Since no human can even come close to solving the problem, it is really the only benchmark we have – even if it is a solution for only 20 users, when 6000 is closer the real problem.

The quick method has two phases. Firstly we must lay down a backbone that ensures that every user can send a message to every other user. We use techniques from a field of mathematics called Graph Theory that let us do this pretty quickly. We end up with a network that connects all users, but still observes the limitations on each computer. This backbone does not need to change very often.

In the second phase, any spare capacity is used to try to connect users that are close in the virtual world with direct or near-direct connections on the network. This boosts their performance even more. This phase is even quicker, and we can re-run it pretty often, to respond to changes in the virtual world.

This technique proved very effective. We generated some random test problems with characteristics like the virtual worlds seen in *World of Warcraft*tm. We were able to deliver all but a few percent of messages in the required time. We have been doing some tweaking on this method, and feel sure that a practical solution can be found that is based on the ideas in this method.

The exact method uses a very mathematically elegant technique that captures all of the elements of the problem in a set of equations. The answer is found by solving the equations. Unfortunately they are not the sort of equation where you can write down the answer straight away. Instead we must solve them step by step, moving closer to the answer each time – the result at one step telling us the direction the next step must be. A bit like groping in the dark in one of those dungeons of old, with nothing but a small (but elegant) candle to guide our way.

Sponsors

University of Wollongong

Host and Financial Sponsor

ANZIAM

Host and Financial Sponsor

Australian Mathematical Sciences Institute

Financial Sponsor

Bluescope Steel Research

Project Sponsor x 2

Geosciences Australia

Project Sponsor

ICT Research

Project Sponsor

Integral Energy

Project Sponsor

Minitab Pty Ltd

Educational Sponsor

The organisers gratefully acknowledge the support of all of the MISG2009 sponsors, without whom the study group would not be possible.