Mathematics and Statistics in Industry Study Group 2007

Industry Project Reports: Equation Free Summaries

Introduction

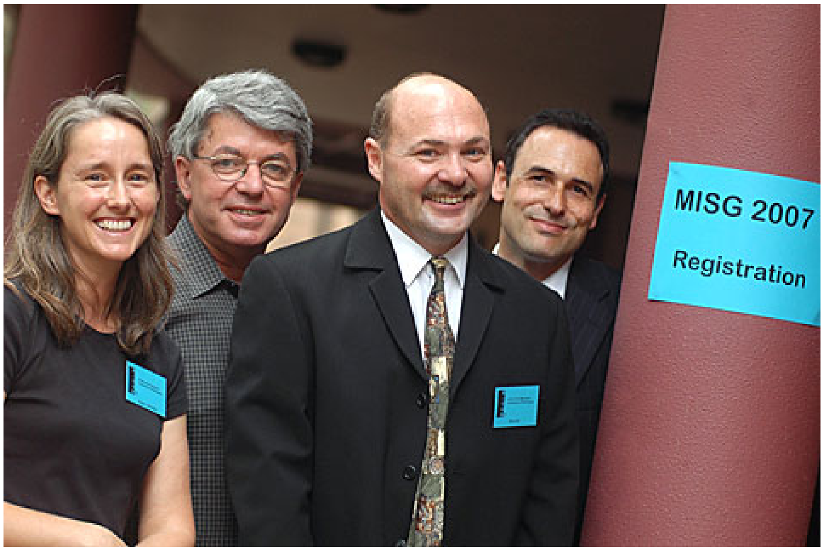

From left: Dr Maureen Edwards - MISG Assoc Director, Prof Robert McKibbin – Massey University, A/Prof Tim Marchant – MISG Director and Mr Joe Maisano - Industry Rep -Trading Technology Aust

The Mathematics and Statistics in Industry Study Group (MISG2007) was held during the week 5-9th February 2007 at the University of Wollongong. MISG2007 was hosted by the School of Mathematics and Applied Statistics at the University of Wollongong and ANZIAM, the Australian and New Zealand Industrial and Applied Mathematics Society. The event was directed by A/Prof. Tim Marchant and Dr. Maureen Edwards. Administrative support was ably supplied by Ms. Joell Hall and Ms. Sue Denny.

Six industry projects were presented; five of these came from Australia and one from New Zealand. Two strong project themes were financial mathematics and electricity supply and generation. About 110 people attended MISG2007 including 20 postgraduate students. In the following pages non-technical outlines of the projects, and the progress made, are presented. Later in the year, the full proceedings of MISG2007 will be published.

MISG2007 was opened by Mr. Stephen Lowe, General Manager Trading, Integral Energy and Prof. Margaret Sheil, DVC-Research, University of Wollongong. The invited speaker was Prof. Robert McKibbin from Massey University. Dr Mike Camden of Statistics NZ and Dr Jeff Dewynne from the Universities of Oxford and Wollongong, gave talks at the student sessions. 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) and the Commonwealth Scientific and Industrial Research Organisation (CSIRO) both supported postgraduate attendance at MISG2007. Thanks go to these organisations and to the University of Wollongong for their financial support.

MISG2008 will be held at the University of Wollongong, 28th Jan - 1st Feb 2008. We hope that MISG2008 will be bigger and better than MISG2007, and hope to see everyone at Wollongong again for next year's event.

Tim Marchant,
Maureen Edwards,
Directors, MISG2007

Index

Strip Track-off and Buckling Between Transport Rolls

Bluescope Steel Research, Wollongong, Australia

Industry Representatives
Dr. Andrew Dixon, Bluescope Steel Research, Wollongong, Australia
Dr. Daniel Yuen, Bluescope Steel Research, Wollongong, Australia
Ms Maria Cozijnsen, Bluescope Steel Research, Wollongong, Australia

Project Moderators
Dr. Barrie Fraser, University of Sydney, Australia
A/Prof. Charlie Macaskill, University of Sydney, Australia
Dr. Mark McGuinness, Victoria University of Wellington , New Zealand

Student moderator
Mr. Aaron Thornton, University of Wollongong, Australia

Bluescope Steel is the leading manufacturer of flat steel products in Australia and New Zealand. Dr Andrew Dixon, at the preliminary presentation, explained that the problem of particular interest was the formation of 'ironed-in' wrinkles in thin steel-sheet products. This is found to occur in a number of environments, but typically is observed in installations where steel sheet is passing over a large number of rollers, some of which drive the sheet and hence provide a tension in the direction of motion. In general, the defects arise at random, and only very infrequently, but can have a significant cost both because of the loss of product and the difficulty associated with tracking down the cause of the problem, which can sometimes be traced to roller misalignment or to metal sheets of non-uniform thickness, for example.

Several different kinds of wrinkles are observed. One type of permanent wrinkle is similar to the folds that are sometimes observed in newspaper printing, and which follow a diagonal pattern arising from the edge of a (usually misaligned) roller. After some early discussion, it was decided that the nature of these wrinkles, and the mechanisms leading to their formation, was relatively well-understood both in the literature and by Bluescope Steel Research. Accordingly, it was agreed that work would concentrate on a second kind of wrinkle that occurs more often in practice in cold steel-sheet rolling and that arises in a fundamentally different way. This second kind of wrinkle usually appears as a single thin ridge close to the centre of the sheet and is more prevalent when convex rollers are used to alleviate mistracking (convex rollers have greater radius near the centreline). The wrinkle may have a longitudinal extent of tens of metres but a width of the order of twenty to thirty millimetres and a depth of less than a millimetre. The ridge typically wanders in the lateral direction around the centre of the sheet.

In order to explain the nature of these wrinkles and how they arise, it was agreed that the participants would concentrate on the following specific objectives:

1. To solve the weakly nonlinear problem for vertical deflection in a rectangular plate under realistic boundary conditions describing the tension imposed by the rollers at each end of the plate ;

2. To understand the mechanisms giving rise to local cross-plate compression, a necessary prerequisite for initial wrinkle formation, in a rectangular plate;

3. To determine the mechanisms for the *iron-on* process as a wrinkle approaches a roller.

Some simple demonstrations with aluminium kitchen foil showed that a thin metal sheet under longitudinal tension and with clamped ends would certainly support small amplitude elastic wrinkles, elongated in the longitudinal direction, giving an appearance like corrugated roofing steel. However, these wrinkles were not permanent and would disappear once the applied tension was reduced. Moreover, such wrinkles, discussed in the literature in a number of contexts, clearly require a locally compressive field as in item 2 above.

In order to study the mechanism for the appearance of wrinkles, the appropriate analytical tool is the Von Kármán equations for the bending of plates, a coupled set of partial differential equations (PDEs) that describe the in-plane stresses and their coupling to the displacement, normal to the plane of the undeformed plate. This vertical deformation need not be small in the full set of equations. If in addition the assumption is made that the vertical deflections are indeed very small, then the equations decouple and the in-plane stresses in the plate obey a (linear) biharmonic equation, so that the stresses can be determined without reference to any knowledge of the vertical deflections. The remaining problem is then to determine the nature of the boundary conditions. The lateral boundaries of the plate are assumed to be stress-free. The ends of the plate, where the rollers are, are assumed to be clamped: this implies that there is tension in both the lengthwise and cross-plate directions.

Under these conditions, the group found that numerical solutions of the biharmonic equation for the stresses reveal a local cross-plate compressive stress quite close to each of the rollers. Furthermore, if a roller is convex, giving rise to a maximum in the *cross-plate tension* at the centreline at the roller location, then there is correspondingly a maximum in *cross-plate compression* at the centreline some distance from the roller. Thus proper consideration of the boundary conditions has provided an explanation for point 2 above.

The above findings explain how a maximum elastic deformation will be found near to the roller location and at the centre of the plate. However, in order to see a permanent distortion of the plate, the elastic stresses must exceed the plastic yield limit, hence ‘setting in’ the distortion. For this to happen the local radius of curvature of the distortion must be quite large: for thin steel sheet we estimate that the vertical deflection divided by the lateral wavelength must be greater than about 1/15. In order to determine this ratio, the full nonlinear Von Kármán equations must be solved in order to determine the vertical deflexion over the plate, in particular at the location of maximum compression. This more challenging part of the problem is the subject of ongoing work.

A further question to be addressed is that although initiation of the ridge-like distortion is apparently random (e.g. a roller is slightly out of alignment), once a distortion is initiated it tends to continue over long lengths of material. One possible explanation for this is that a local ridge in the plate when reaching the roller will lead to a modified tension at the roller that in turn influences the compression upstream: this proposed effect can only be explored with careful solution of the full governing equations with small perturbations in the roller boundary conditions.

Finally, the group made various approximate calculations, related to the *ironing-in* of a buckle, to make it a permanent wrinkle at the roller. Critical contact lengths between the steel sheet and the roller, such that friction will hold a buckle in place while it is being compressed, were found to be less than a millimetre. The coiling pressure was calculated, and was found to be enough to compress a buckle down to a size of the order of 50mm across, which compares well with data provided by Bluescope Steel Research.

In conclusion, the following progress was made over the week of the MISG meeting. Firstly, various mechanisms were identified for the formation of wrinkles in thin steel sheet. Secondly, once attention was focussed on the type of wrinkle that occurs as a central ridge, mathematical and computational analysis indicated that by proper choice of boundary conditions at the roller, a local compressive in-plane stress could be produced that would provide a localised maximum distortion in the sheet. Approximate calculations confirm that there is enough tension in the steel to make the distortion a permanent wrinkle when it passes around a roller.

With this information, it has been possible to hypothesise the mechanisms by which a central wrinkle is created by plastic distortion and then set in at the roller. Finally, it has become apparent that to confirm this hypothesis the full Von Kármán plate bending equations will have to be solved in order to determine the maximum vertical deflection. This is the subject of ongoing work.

Application of the Non-stationary Travelling Salesman Problem to Maritime Surveillance

Air Operations Division, DSTO, Fishermans Bend, Melbourne, Australia

Industry Representatives
Dr David Marlow, Air Operations Division, DSTO, Australia
Dr Jason Looker, Air Operations Division, DSTO, Australia

Project Moderators
Dr Philip Kilby, National ICT Australia, Australia
Dr Patrick Tobin, Swinburne University, Australia

Student moderator
Ms Ruth Luscombe, University of Melbourne, Australia

The survey and identification of ships within and near to Australian waters is an important national security issue. With limited reconnaissance capability and a huge area to survey it is very important that this surveillance is done as efficiently as possible. Two important factors arise in the survey. The number of ships which escape detection and/or identification must be minimized and the flight path of the detection aircraft needs to be as short as possible.

The layout of ships on the ocean forms a network of nodes. In a classic Traveling Salesman Problem (TSP), the nodes in a network need to be traversed by as short a path as possible which visits each node. Such problems have traditionally been attacked with heuristic procedures like simulated annealing and genetic algorithms which cope with the rapid increase in size of the problem which comes with increasing node numbers. In a classic problem the nodes are stationary but surveillance involves moving target nodes in the form of ships and there are other variable factors to consider such as the identification horizon the range at which the aircraft can identify the ship involved.

In this work we examine a non-stationary TSP. A survey aircraft follows a path through a specific allocated zone known as the Area of Interest (AI). Within this zone certain specific points on the flight path known as waypoints are laid down, and the survey aircraft follows a path which winds through the AI progressively detecting ships by radar in a detection circle up to 100 nautical miles in radius, and proceeding to identify the ships by flying to within an identification region, possibly as far as 20 nautical miles from the ship and possibly at the ship itself. The survey aircraft must visit the waypoints in a specific order. This is an important constraint in the solution. The ships may move with speeds and directions which can be found upon detection or they may be stationary, as may arise in fishing for example.

Various types of aircraft are used in surveillance from helicopters which cruise at 100 knots to jets which may travel at up to 350 knots. Other aircraft with speeds up to 450 knots are available to purchase, but are not currently in use. Various survey region configurations may also be used and their sizes varied, possibly to suit the aircraft used and region location.

DSTO were seeking advice about a number of aspects of the surveillance. In particular they were interested in how important it is to regard the ships as non-stationary. It was of concern to know whether high relative speeds by aircraft over ships would render this aspect unnecessary. Since ship speeds can be about 30 knots the speed could be very significant for a 100 knot survey and rather less so for a 350 knot survey!

In practice ships will move in and out of the AI. One important rule laid down in surveillance is that the survey aircraft must stay in the AI until it leaves its mission. This constraint means that some ships will be detected at the fringe but may escape identification if they leave the AI before the survey aircraft can reach them. Some ships will escape detection altogether as they pass through the region beyond detection limits.

The work described here was done on a standard AI which was a square of side 300 nautical miles. In the first instance ships were left stationary, the detection limit was set at 100 nautical miles and the identification limit was set at the ship location. Simulation runs on a genetic algorithm in Matlab suggested that the total path generated would be of the order of 2.7 times the length of the straight line path linking the waypoints a path we may call the wayline here. Further simulation runs suggest that this path may be reduced by up to 25% in length if the identification range is extended to 20 nautical miles. This confirms that using a zero distance identification range will give a pessimistic forecast.

It is also clear that a TSP solution will generate sharp cornered paths which cannot be sustained in real terms because of acceleration effects on the pilots and aircraft so that some turning circle will also need to be incorporated and this can lengthen the path. This is still an issue if unmanned reconnaissance is used.

It was of interest to assess the actual ship detection efficiency. Various simulation runs on a single leg of the path using a more restricted detection range of 50 nautical miles suggest that detection efficiency is up to about 90% but this will decline with increasing ship speed. This will improve if the detection range is increased.

In the model generated, a program progressively and repeatedly calls on implementing a simple TSP algorithm known as 2-Opt. This is more efficient than the genetic algorithm currently used and has very rapid solution times for the number of nodes generally called upon within the detection radius. Ships were laid down randomly on a large square region nine times the size of the centrally embedded AI and simulated with a specified speed and randomly allocated directions of motion.

A simulated craft traverses the region calling at waypoints in order and using the next waypoint as the final point of the currently operative TSP. This enforces the waypoint travel order condition. As the craft traverses the region, the detection region changes and various triggers are used to generate new TSPs. These include proximity to the wayline for detected ships and new ships coming to the detection horizon. Identifying which ships are not worth chasing and whether the current path is over budget in terms of time (and hence effectively fuel) become key aspects of the solution. Examining the path progressively with a moving standard deviation or similar statistical procedure may assist in keeping the survey to the budget constraint.

Repeated simulations suggest that it is significant that ships move and that no current speed ratio of ships to aircraft should ignore this point. The simulations work on minimizing flight path by implementing a TSP repeatedly but also generate identification efficiency as an outcome. These give a guide to the survey group on what parameter combination options will be acceptable. For example the AI size and the aircraft used are parameters in the control of the survey group, whereas ship speeds and directions of travel and actual ship numbers are out of their control. Features like the detection and identification limits may be subject to weather variations.

An exhaustive preliminary analysis of surveying a region with ships of all levels of speed and various other parameter options may be done in advance to give aircraft a better idea on which ships are worth chasing at all and which may be better left to attempt to reach later in the reconnaissance.

One standard algorithm which repeatedly calls on an efficient TSP subroutine is sufficient for all cases of ship patterns. Overall the current work has substantially extended the use of non-stationary TSPs in this field and shows promise for substantial future development. Several issues remain unresolved. These include the best response to the case where no ships are in detection - a likely outcome where ships are sparse. In this case, the best procedure for the aircraft to return towards the wayline needs investigation. Another issue is the extent and cost of backtracking to pick up stray ships. This action needs to be quantified in a cost-benefit analysis if it can increase risk on meeting budget.

Calibrating Mean Reverting Jump Diffusion Models

Integral Energy, Sydney, Australia

Industry Representative
Dr Andrew Ziogas, Integral Energy, Australia

Problem Moderators
Dr Jamie Alcock, UQ Business School, The University of Queensland, Australia
A/ Prof. Ramaprasad Bhar, School of Banking and Finance, UNSW, Australia

Student Moderator
Mr Tony Vassallo, The University of Sydney, Australia

In recent years the Australian National Electricity Market (the NEM) has been undergoing the transition to a fully deregulated marketplace. In December 1998, for the first time, the wholesale price of electricity was subject to market forces. The NEM includes the Queensland, New South Wales, Victorian, South Australian, Tasmanian and Australian Capital Territory electricity markets. Almost $8 million end users are supplied by the world's longest interconnected power system, and the NEM trades up to $7 billion of electricity annually [5].

NEMMCO Ltd (The National Electricity Market Management Company) was established in 1996 to manage the NEM, a role which carries with it the responsibility for setting the spot price. The spot price is determined via a sellers dutch auction. Each day, each generator submits a complex bid of prices and volumes. The demand fluctuates throughout the day, and every 5 minutes short-term supply and demand are realigned by NEMMCO. The 5 minute dispatch price for all bidders is set to the winning bid of the marginal supplier, and six sequential dispatch prices are averaged to determine the half-hour spot price. All successfully bidding generators receive the spot price for their product. Currently spot prices are artificially bound by NEMMCO to remain between -$1,000 and $10,000 per MWhr.

The introduction of market forces to the NEM has provided a plethora of challenges for mathematicians, economists and financial economists[1] amongst others. As electricity is not storable the spot price process is extremely volatile. However the trend underlying the spot price process is highly predictable and highly periodic. The supervolatility of spot price stems from a combination of unexpected events. Unexpectedly high demand (possibly due to an unexpected temperature change) or unexpectedly low supply (due to an unscheduled generator outage or distributional failures /constraints) results in a rapid increase

[1] For example see recent and significant works by [2,3,6]

in price. There is a strong positive relationship between the spot price spikes and large changes in demand (load).

Retail customers typically pay a fixed price for electricity. Hence the electricity retailer manages a portfolio of floating-for-fixed swap instruments. The retailer receives a fixed payment for electricity and pays a floating rate to the generator. Managing financial risk in such a volatile market is a formidable challenge. The magnitude of this challenge is highlighted when one considers that the retailers financial risk is the product of load and spot price.

One model that has been heavily utilised in the literature is the Mean-Reversion Jump Diffusion (MRJD) model. The clear relationship between load and spot price was investigated in order to improve the predictive ability of the raw MRJD model. Several key characteristics of the data were identified:

1. The expected spot price is a linear function of the observed load.

2. The size of the spot price spikes is a linear function of several variables, including the previous spot price. The size of the spot price spikes is also dependant upon whether or not it is winter. However, the marginal effect of these variables is small. By far the most important predictor of spot price spike size is change in observed load. Curiously the price sensitivity to load is higher in winter than in summer.

3. No spot price spikes occur in either Autumn or Spring, nor do they occur during the night (between 8pm and 8am).

4. The volatility of spot price is a periodic function. However load is also highly periodic, so it is assumed that volatility of spot price is proportional to load.

Each of these characteristics is incorporated in the MRJD model to be calibrated. We examine several potential methods of calibration of general MRJD models, with the objective of calibrating the model specified in the previous section. Maximum Likelihood Estimation (MLE) has traditionally had difficulty with calibration of MRJD models, due primarily to its inability to distinguish between a small jump and a large Wiener movement. Indeed, this difficulty is common to many methods when calibrating MRJD models. We examined two alternative calibration methods, namely the Generalised Method of Moments [4] and a simulation based estimation procedure [1].

1. J. Alcock and K.Burrage. A genetic estimation algorithm for parameters of stochastic ordinary differential equations. *Computational Statistics and Data Analysis*, 47(2):255--275, 2004.

2. K. Chan. *Modelling Short-Term Interest Rates and Electricity Spot Prices*. Ph D Thesis, UQ Business School, The University of Queensland, 2006.

3. L. Clewlow and C. Strickland. *Energy Derivatives: Pricing and Risk Management*. Lacima Publications, 2000.

4. L. Hansen. *Large sample properties of generalized method of moments estimators*. Econometrica, 50(4):1029--1054, 1982.

5. NEMMCO Ltd, Melbourne. *An Introduction to Australia's National Electricity Market*, 2004.

6. P. Simshauser. *Microeconomic Reform of Wholesale Power Markets: A Dynamic Partial Equilibrium Analysis of Negative Externalities*. PhD thesis, The University of Queensland, 2001.

Calculation of a Risk Measure for the Net System Load Profile

Integral Energy, Sydney Australia

Industry Representatives
Mr Alan Prictor, Integral Energy, Sydney, Australia
Dr Andrew Ziogas, Integral Energy, Sydney, Australia
Mr Justin Arundell, Integral Energy, Sydney, Australia

Moderators
A/ Prof. Ken Russell, University of Wollongong, Australia
A/ Prof. Pietro Cerone, Victoria University, Australia

Student Moderator
Ms Vivien Challis, University of Queensland, Australia

Integral Energy is one of three franchises which provide retail electricity in New South Wales (NSW). Integral Energy purchases wholesale electricity from the National Electricity Market and sells this to retail customers. The electricity market is unusual because the price at which electricity is sold to retail customers is fixed but the price that the electricity retailer must pay for electricity from the National Electricity Market changes. Thus price changes incurred by Integral Energy are not passed on to their customers, introducing electricity price risk. Integral Energy uses the electricity hedge market, customer contract management, and the Electricity Tariff Equalisation Fund, along with market forecasting, to reduce its exposure to electricity price risk. These contracts are used to hedge Integral Energy against risk for electricity demand within the 95% confidence intervals of their long range electricity demand forecasts.

This project is in part concerned with comparing Integral Energy's electricity price risk with that of the other two NSW electricity franchises: Energy Australia and Country Energy. Integral Energy want a measure of the relative volatility of their demand compared with that of the other two franchises. Integral Energy is most at risk during very high load periods, such as very hot days in summer and very cold days in winter. Integral Energy's customers are largely based in Western Sydney, where it gets hotter in summer and cooler in winter than on the coast. Houses in this area also tend to be less energy efficient, and Integral Energy believe these factors combine to produce higher demand during peak periods. Energy Australia's customers are largely based on the coast around Sydney. Coastal temperatures are more stable than inland, which Integral Energy suggest results in less volatile demand. Country Energy largely provides electricity for the rest of the state and from historical data appears to have the least volatile demand.

Integral Energy provided data collected every 30 minutes from 1 January 2002 to 31 December 2006. The actual measurement values were given as well as Integral Energy's long term forecast values and their 95% confidence intervals. The measurements included temperature at Bankstown (in western Sydney), the National Electricity Market electricity price, and various load measurements. For this project Integral Energy wanted us to consider the Net System Load Profiles (NSLP) for each franchise. The NSLP are a measure of the load over which the retailer has no control; it is discretionary electricity consumption controlled by the customers. The NSW State Load was also found to be of use in the modelling.

The brief from Integral Energy was to find a model for their NSLP given the temperature in Bankstown, which is used as a proxy in lieu of detailed temperature variation for the francise, to estimate discretionary load. This model would in effect be used to calculate expected load based around temperature in the short term. The second major task was to quantify the difference in volatility between the three franchise NSLP profiles.

The problem was approached by first undertaking an exploratory data analysis of the 87,648 data points. The investigation of the profiles led to a realisation that predictors other than temperature would be necessary in the model to reliably predict Integral Energy's NSLP. Predictors which the data analysis suggested should be included were: current Bankstown temperature; type of day (namely, working or non-working); time of day; month of the year; year; and, an interaction between month and time of the day. This interaction term allows for days having differently shaped NSLP profiles at different times of the year. There were very high correlations between current Bankstown temperature and the Bankstown temperature 24 hours earlier, and similarly between current Integral Energy NSLP and that 24 hours earlier. It was confirmed with Integral Energy that a model predicting load 24 hours in advance would be useful to them. The temperature at the time load is to be predicted and the temperature 24 hours in advance were chosen as predictors for NSLP. This is based on the assumption that Integral Energy will be able to get very accurate Bankstown temperature forecasts from the Australian Bureau of Meteorology. However, since NSLP are not known until at least three weeks after the time in question, using NSLP 24 hours earlier would not be a practical predictor to have in the model. The NSW State Load 24 hours earlier was used as a proxy for this in the model, as it is a quantity known to Integral Energy in real time. As noted earlier, Integral Energy is particularly at risk during peak load periods and therefore accurately predicting the NSLP at these times is important. An increment when temperatures are over 28^{o} C was used in the model to aid in accurately modelling the NSLP in peak periods.

The model described above was fitted using the NSLP for Integral Energy for 2002 - 2005 inclusive, retaining the 2006 data for validation purposes. The model described 86% of the variation in Integral Energy's NSLP. With such a large volume of data this value is adequate. A value of 0.165 was obtained from the model for the root mean squared error which measures the variability of the errors. The model explained 83% of the variation in the data when predicting the NSLP for Integral Energy for 2006. Although these results are promising, a closer examination of the errors showed that the model was still unable to predict the load accurately during peak periods.

Models with the same terms involved were fitted to the Country Energy and Energy Australia NSLP and the models explained 88.5% and 88.6% of the variation respectively. These values were 86.0% and 89.0% respectively when the models were used for predicting the 2006 NSLP. The root mean squared error was 0.1363 for Energy Australia and 0.1079 for Country Energy. Compared with the Integral Energy root mean squared error above, these do suggest a higher volatility of NSLP for Integral Energy than for the other two franchises. However, the data we were given was normalised to a mean of unity, and so Integral Energy will need to re-scale these measures of variability in order to get the correct relative variability.

Work is continuing to improve the prediction of the model during peak load periods, and the errors will be modelled in order to provide better measures of the variability of the NSLP for each franchise.

From consideration of the problem and the data, other predictors may prove useful to Integral Energy in modelling the NSLP load. In particular, using temperature measurements from around the franchise may provide more insight than just using the Bankstown temperature. Different observations may also be weighted (given more importance) on the basis of customer or usage density. Humidity may be a key predictor, as suggested in Doulai & Cahill (2001). Also, other socioeconomic factors and demographic factors may be useful in predicting NSLP more accurately.

1. Doulai, P. & Cahill, W. (2001). *Short-term price forecasting in electric energy market*. Proceedings of the International Power Engineering Conference (IPEC2001), Singapore, May 17-19, 2001, 749-754.

Operating and planning an electricity transmission grid to maximize the contribution of wind

Transpower Ltd, Wellington, New Zealand
Energy Efficiency Conservation Authority, Wellington, New Zealand

Industry Representatives
Mr Conrad Edwards, Transpower Ltd., New Zealand
Mr David Rohan, Energy Efficiency and Conservation Authority, New Zealand

Problem Moderators
Dr Geoff Pritchard, University of Auckland, New Zealand
Dr Winston L. Sweatman, Massey University, New Zealand

Student Moderator
Mr Kim Nan, University of Auckland, New Zealand

The power generated from wind energy in New Zealand could significantly increase in the future. The focus of the MISG problem is how to design and operate the national power grid to facilitate this. In character, wind power generation is rapidly fluctuating and unpredictable. Two ways in which this variability may be managed are wind/hydro matching in which hydro stations are used to compensate for drops in wind power, and wind sloshing in which by having a number of wind farms the 'random' fluctuations in one farm are partially offset by changes across the farms collectively. In both cases the power balance can only be maintained if there is sufficient headroom in the transmission lines. So an optimal value for this headroom needs to be found.

The MISG group generated a number of simple models to investigate different aspects of the problem. They started their study with a model incorporating just three power stations. At one node of the system a fast-ramping wind farm and a slow-ramping cheap power station were situated. These were connected by a transmission line to a fast-ramping expensive power generator which was next to the power load on the system. As the wind power generated rapidly varies, the power output of the cheap power station is adjusted to try to balance the load. As this cheap power station is slow-ramping it cannot immediately compensate for rapid changes in the wind power and the balance must be met by the fast-ramping expensive power generator (which could be a hydro power station). This model enabled the study of some aspects of how best to balance the sources of power generation for the cheapest long term operation.

Within this first model an exact analysis can be made by assuming the wind-power to be a square-wave or a similar simplified function and the load to be constant. For this the mathematical solution typically involves finding the minimum point of a U-shaped quadratic function. Minor adjustments to this toy problem provide further insight.

When wind-power suddenly drops it takes time to increase the slow-ramping cheap generator to meet the power demand. The deficit must be met by the expensive power station. Even if unlimited wind power were to be available there is usually a point at which it is more effective to balance wind power with a contribution from the cheap power station in order to constrain the amount of power generation required by the expensive power station when the wind drops.

An alternative possibility illustrated by this simple model occurs if an increase in wind-power generation can be anticipated. In this case it may be worthwhile to deliberately increase power generation from the fast-ramping expensive station and ramp down the slow-ramping cheap station prior to the onset of the increased wind-power generation. This enables the wind to be more fully used as soon as it is available.

Beyond the toy problem members of the group considered more realistic versions of this simple model. Theoretical approaches were generated for tackling the issue of finding the optimum strategy.

The three-power-station model was also simulated on the computer. For this, more realistic models of the wind and wind generation were used. The program did not anticipate wind changes: it only increased and decreased the cheap slow-ramping power generator in response to the wind. For each simulation a constant minimum level of cheap power generation was maintained. The remaining load provides a headroom which can be filled by wind-power generation. If the wind power generated is lower than this headroom then the slow-ramping cheap power generation can be increased to compensate, the difference meantime being met by the expensive generator. If, instead, the wind-power generation potential is higher than the headroom, then wind is spilled.

The simulations were made with headrooms from 1% to 100% of the load. For each headroom the cost of meeting the power load demand was found for 100 random realisations of the wind time-series and the average of these was found. Again the relationship between cost and headroom was a U-shaped curve with an optimal headroom. Simulations of this kind could be readily extended and adjusted.

Further models allowing for additional generators in more complicated networks were also studied. One way to approach these problems is to assume that the power generation from the slow-ramping power stations is kept constant. If fluctuations in wind-power generation are very rapid relative to the slow ramping in these stations, then they cannot be expected to meet any significant portion of the fluctuating demands.

A first model takes the three power generators, as in the previous simple model, but adds a further slow-ramping medium-priced station alongside the expensive fast-ramping power station at the point of demand (the cheap and wind power generators are again separated from these generators by a transmission line). This model highlights some of the constraints due to transmission line capacity. A balance needs to be struck between the use of the two slow-ramping power generators, in order to enable wind power to be passed along the transmission line when it is available. By using the medium-priced generator, instead of the cheap generator, headroom is created within the transmission line. As before, in the examples studied there was an optimal solution.

Clearly the time profile of the wind and its predictability are important in problems of wind generation. Members of the MISG team studied aspects of a set of real windspeed data. In order to use windspeed data for wind-generation planning it must be converted into its potential for power-generation. As well as studying this particular dataset a two wind-farm network was simulated. Further possibilities were identified for existing windspeed data.

Overall, the MISG group studied the Transpower/EECA problem using a number of approaches and simple models. Some of these approaches will be suitable for further extension. The work has highlighted the fact that, although wind-power generation provides great opportunities for meeting energy requirements, the planning and management of its transmission presents new challenges. These will have to be met in order to fully utilise this resource.

Determining the independence of various measures of financial risk

Trading Technology Australia, Sydney, Australia

Industry Representatives
Dr Alex Radchik, TTA, Sydney, Australia
Mr Joe Maisano, TTA, Sydney, Australia

Moderators
Prof Ed Maberly, Monash University, Australia
Dr Andy Wilkins, University of Queensland, Australia

Trading Technology Australia (TTA) was founded in 1996 and has a background and client base in the financial markets. In 2002, TTA delivered its first Energy Markets project in the form of a risk management report to the CEO of a large electricity retailer. Since then, TTA has been working to further understand energy market issues such as the relevance of particular models used in trading and risk.

Common measures of financial risk include Value-at-Risk (VaR) and Earnings-at-Risk (EaR). VaR measures the probable change of a portfolio's position due to market movements within a given confidence interval. EaR measures the probable loss in earnings due to market or volumetric movements within a given confidence level.

The question posed by TTA to MISG2007 was: How independent are VaR and EaR? To answer this question, the MISG2007 team focused on a simple example portfolios for an electricity retailer, consisting of forward contracts with the generators. Given the forward price curves, the cumulative probability distribution for portfolio's position (relative to the present position) at some future time may be formed, and the 99% VaR is the value for which this is 1%. The retailer must sell electricity to consumers at a fixed price, but buy electricity from the generators at the spot price or the forward price. Given the future predictions for the prices and the demand, the cumulative probability distribution for the retailer's cash flow at some future time relative to the present cash flow may be formed, and the 99% EaR is the value for which this is 1%.

To enable risk reporting, the electricity retailer uses a bootstrap approach on relevant historical data to compute probability distributions for the spot, forward prices and electricity demand in the future. These are computationally expensive to calculate, and if a map from VaR to EaR could be found, significant savings would be generated for the retailer by making their risk reporting framework more streamlined, transparent and risk-compliant.

Both risk measures increase with market volatility, but accurately quantifying EaR given VaR does not seem to be possible. Two factors spoil any chance of useful relationships between VaR and EaR. The first is that the portfolio's value (used to calculate VaR) depends on the forward price curves, while the predicted cash flow (used to calculate EaR), depends on the future spot price and demand. In some markets the forward prices are strongly related to the future spot prices, but for electricity markets there is little correlation. Secondly, VaR is sensitive to drops in prices, while EaR is sensitive to increases in prices and demands. Mathematically, VaR is looking at left tails of probability distribution functions, while EaR is looking at right tails. The generic case is asymmetric distributions, so there are no generic maps between VaR and EaR.

Computational time for calculating the various (joint) probability distributions may, however, be reduced in the following cases. (1) The forward prices are strongly correlated or equal to the spot prices. (2) In electricity markets, the spot price is set by the demand and the generators' bid prices into the pool. While the latter is random, it may have relatively small volatility, meaning that the set demand, bid price) might be more efficient to work with than (demand, spot price). (3) The probability distributions would be simple to calculate if the underlying stochastic processes could be found and calibrated from historical data. In this method, it is the calibration (and continual recalibration required by regulators) that is difficult.

Sponsors

	Australia New Zealand Institute of Applied Mathematics Financial Sponsor
	University of Wollongong Financial Sponsor and Host of MISG 2007
	Australian Mathematical Sciences Institute Financial Sponsor
	CSIRO Mathematical and Information Sciences, supporting student attendance at MISG 2007 Financial Sponsor
	Bluescope Steel Project Sponsor
	DSTO Project Sponsor

	Energy Efficiency and Conservation Authority Project Sponsor
	Integral Energy Ltd Project Sponsor
	Trading Technology Australia Project Sponsor
	Transpower Ltd Project Sponsor
	Ceanet (Maple 11) Educational Sponsor
	Thomson Educational Sponsor