



# Mathematics-in- Industry Study Group 2004

*Problem reports: Equation-Free  
Summaries*





Centre for Mathematics in Industry

## **STRIP TEMPERATURE IN A METAL COATING LINE ANNEALING FURNACE**

**New Zealand Steel**

**Industry Representatives:**

Philip Bagshaw, Manager Product Technology, New Zealand Steel  
Nebojsa Joveljic, Principal Technologist, New Zealand Steel  
Damien Jinks, Research Scientist, Bluescope Steel Research, Australia  
Michael O'Connor, Senior Process Control Engineer, New Zealand Steel

**Moderators:**

Mark McGuinness, Victoria University of Wellington  
Stephen Taylor, The University of Auckland

New Zealand Steel (NZS) use a unique process to convert New Zealand iron-sand into steel sheet products at its Glenbrook mill near Auckland. Traditional galvanised steel (Galvsteel<sup>TM</sup>) and the new product Zincalume<sup>^</sup> are produced in a range of dimensions, grades and coating weights.

The steel strip is annealed prior to being coated, by heating to a predetermined temperature for a definite time. Annealing produces desirable changes in the crystalline structure of the steel, allowing NZS to tailor its strength and ductility. Strips of steel sheet are passed through a 150m long, 4.6 MW electric radiant furnace at speeds of up to 130 metres per minute in order to achieve the strip temperatures required for annealing, and subsequent coating. The temperature along the furnace is controlled by varying the power supplied to the heating elements and by use of cooling tubes. The cooling tubes are located in the last half of the furnace and consist of steel tubes through which ambient air is pumped. It is important that steel exit the furnace with the correct temperature for the coating that is applied at the exit point.

The line speed through the furnace is reduced for strips of large thickness and width in order to achieve the required temperatures. At the beginning of the annealing--coating line there is an automatic welding process which welds the beginning of a new coil of steel sheet to the end of its predecessor, allowing the line to run continuously.

In each of the twenty zones of the furnace, there are thermocouples in steel tubes, which are used to measure furnace temperature. The thermocouple temperatures are compared with desired temperature set-points, and the heating elements are controlled accordingly. Steel strip temperature is also measured, using non-contact pyrometers at three positions in the furnace.

If there is no variation in strip dimensions and annealing settings then the line is able to run in a *steady state*, with the furnace temperatures remaining steady at the desired thermocouple settings. NZS have already developed a mathematical model of furnace and strip temperatures for this steady state operation. Challenges occur

when there is variation in strip dimensions or annealing settings because the furnace-strip system has a large amount of thermal inertia. Consequently the line is in a *transient* state for up to 50% of its operation, with varying effects on quality control of the product.

Two improvements are planned for the line in the very near future; a 3 MW induction heater and a gas jet cooler. The induction heater is capable of heating the strip rapidly. The steel strip will pass directly from the induction heater into the radiant furnace. The extra heating power should allow the system to achieve greater line speeds for strips of large thickness and width. Further, with its more rapid response, the induction heater has the potential to reduce the time spent in transient modes of operation. In the gas jet cooler, which will replace part of the existing cooling zone, cooled furnace gas is blown directly onto the steel strip. The new cooler section is expected to respond more rapidly than the existing cooling tubes, giving more precise control of dipping temperatures.

NZS set the following tasks for the Study Group:

- Develop a mathematical model for transient furnace conditions.
- Investigate the accuracy of the existing steady state model.
- Predict transient strip temperatures for actual production schedules with changes in product dimension, steel grade and furnace temperature settings.
- Couple the temperature model to a metallurgical model.

The Group focussed first on modelling the temperature of the steel strip, and soon discovered that this can be accurately modelled as a function of time and just one spatial coordinate, the distance from the entry point of the furnace. Temperatures rapidly equilibrate across the thickness of the steel. Thermal diffusion along the strip was found to be negligible for the length of time that any part of the strip was in the furnace.

This strip model led to a steady state model for the furnace--strip system which took into account the power supplied by heating elements. The Group then compared this to NZS's steady state model. The models differ in that MISG's model allows for continuous changes in temperature along the length of the furnace while NZS's model is discrete, involving one value of strip temperature and one value of the furnace (wall) temperature for each of the furnace's twenty zones. Calculations indicated that the models were in good agreement, confirming the accuracy of the NZS steady-state model.

Next the Group studied the heating of the furnace walls and calculated that while it would take hundreds of hours for bulk changes to wall temperatures, the inner surfaces of the walls heat up rather quickly. They respond to radiation changes on a time scale of about one minute, and to a depth of only a few millimetres. This is too rapid to account for the observation that the furnace typically takes five minutes to equilibrate, so attention was then shifted to the steel hearth rolls (the rollers which carry the strip along the furnace), to see if they could be the main source of thermal inertia within the furnace. Preliminary calculations indicated that the hearth rolls do indeed respond to temperature changes on the correct time scale. A transient model for steel strip temperature was developed that included the hearth rolls.

During the meeting the Group realised the importance of *view factors*. These factors need to be calculated to accurately model the radiation exchange between the steel strip, heating elements, furnace walls, cooling tubes, hearth rolls and thermocouple tubes. This will improve the accuracy of the transient model for the temperature of the steel strip, by giving a more accurate understanding of the radiation environment that the steel strip is experiencing from moment to moment.

The thermocouple tubes house the thermocouples which are used to estimate the furnace temperature in each zone of the furnace. They play a vital role, as they are used to control the power fed to the heating elements. The temperature measured by the thermocouples is not the true furnace temperature, which is needed for accurate

calculation of steel strip temperature, in both the transient and the steady-state models. A calculation of view factors for these tubes will give much better information on the true furnace temperature than is presently available.

## **MODELLING OF A POULTRY SHED**

**NRM/Tegel Ltd**

### **Industry Representatives:**

John Foulds, NRM/Tegel Ltd, Auckland, NZ  
Robert Lloyd, NRM/Tegel Ltd, New Plymouth, NZ

### **Project Moderators:**

Robert McKibbin, Massey University at Albany, Auckland, NZ  
Andy Wilkins, Canesis Network Ltd, Brisbane, Qld, Australia

Tegel Foods is New Zealand's leading producer and supplier of poultry products, providing an extensive range of quality poultry products to New Zealanders for over thirty years. Tegel is part of the Heinz-Wattie group of companies, owned by multi-national food producer HJ Heinz Co. Tegel Foods began operations as a department of General Foods Corporation in 1966 and now employs approximately 1700 people at its sites throughout New Zealand.

Tegel Foods is a fully integrated poultry producer involved in breeding, hatching, feeding, growing, processing and marketing of chicken and turkey in New Zealand. Its product range includes fresh, frozen and cooked whole chickens and fresh and frozen chicken portions. NRM New Zealand markets all feed and animal health products that are sold externally.

The problem presented by Tegel was specifically to model the energy exchange between the chickens and their shed environment in order to better understand and control the shed climate and thereby maximize growth rate.

A typical shed has chickens placed as day-old chicks at a stocking density of about 21 birds per square metre. They are reared on a concrete floor (about 15 cm thick), with a 5 cm layer of dry wood shavings spread as "litter". This litter remains with flock for the duration of the batch, "composting" down to a friable litter material consistent with "50% sawdust mixed with 50% dry garden soil". The sheds are of "controlled environment" type, and the birds are grown within a specific temperature profile as they get older. The shed temperature control starts at 32 °C at the day of placement, reducing down about 0.4° per day to 20°C by the time the birds reach final processing age (average 37 days). The chickens have unlimited access to feed

and water, and grow to a specific growth profile with target weight-for-age expectations. Specific air exchange requirements are necessary to maintain a shed environment acceptable for animal welfare and performance parameters. Water generated into vapour/humidity, through evaporation, and CO<sub>2</sub> are the predominant waste products which must be removed.

The moisture content of the dry wood shavings prior to placing the chicks is close to 5%. By the end of the growing cycle the litter moisture is ideally no more than 20%. Water accumulation in the litter is insignificant compared to total water throughput during the run. The air exchange is determined by total biomass within the shed and therefore increases throughout the life of the flock. Failure to remove sufficient waste air leads to "wet litter" which causes welfare problems as well as performance depression expressed by low feed intakes, low weight gains and poorer feed conversion.

As the birds grow, progressively generating their own body heat, the supplementary heat requirement in the shed decreases and the need to remove heat from the shed starts to overlap. This transition from a heating to a cooling mode is strongly influenced by the weather conditions outside the shed, combined with insulation values of the shed, weight for the age of the flock and target shed environment temperature.

These daily shed temperature targets are based on achieving the optimum "comfort" of the birds at every stage. However as the biomass increases and the influence of heat build-up occurs at floor level, then cooling requirements become harder to formulate on a mathematical basis. Daily temperature monitoring normally measures "ambient" air temperature 30cm above the birds' heads. This temperature is therefore not an accurate temperature requirement but an assumption based on visual flock behaviour. This temperature "perceived" by the birds is a combination of ambient shed temperature, relative humidity, air flow, metabolic heat production and litter temperature.

The modelling of the shed environment's inputs and outputs will be particularly valuable for assessing three fundamental inputs of economic importance: feed nutrient density in terms of energy formulation, heating in terms of gas/power usage, and heat removal via extraction fans. Optimisation of liveweight gain and feed conversion potential are the end targets.

The important variables needing careful consideration include:

- \* The temperature and relative humidity outside the shed.
- \* Supplementary heating into the shed.
- \* Energy and nutrient density of feed consumed by the chickens.
- \* Increase in biomass within the shed.
- \* Heat accumulation and "storage" in the litter and floor under the chickens from biomass heat generation.

- \* Heat generated by composting effect of litter bed.
- \* Increase in insulating effect of birds on litter heat from increasing biomass.
- \* Effect of air flow on heat transfer.

\* \* \* \* \*

The MISG group found that Tegel's farmers raise their chickens in sheds of rough size 15 wide, 80 m long and 3.5 m high. Between 30,000 and 40,000 one-day-old chickens are introduced to the shed where they are kept with unrestricted access to food and water for between 30 and 40 days, at which time they are between 2 and 3 kg in weight.

The shed's floor is concrete. On this is the litter, which is initially wood shavings which then gets mixed with chicken manure, a good deal of which is excreted water which must be removed by ventilation. The shed's ceiling and walls are well insulated. When the chicks are under about 2 weeks old the shed is heated to between 30 and 35 °C, with minimal ventilation. After that time the chickens are weaned off the heat, and the shed may be intensively ventilated, depending on the interior and exterior climatic conditions.

A field trip to one of Tegel's sheds convinced us that the shed could be treated fairly accurately as a homogeneous structure; the air seemed to be well-mixed and the chickens and litter were spread evenly across the shed floor. Therefore the model of the situation consisted of stratified layers: at the bottom was the soil below the shed, then the concrete and then the litter; above that was the "chicken layer", then the internal air, the shell of the shed, and finally the external air.

Our model of the situation included the temperature of the external air, the internal air, the chickens, the litter, the concrete floor and the underlying soil. It also included the relative humidity (RH) of the external air, the internal air and the moisture content of the litter.

The main input of heat into the system was through the metabolism of the chickens. Experimental data suggested that a sufficiently precise model for a chicken was that its heat and moisture output was proportional to the surface area of its lung. A typical 2 kg bird produces roughly 10 W of sensible heat, and respire 0.28 kg of water per day. For a shed of 30,000 birds, this is about 8 tonnes/day, or 0.1 kg/s.

The heat from the chicken passes into the air, but is also used to evaporate moisture from the litter. Much of this is usually vaporised because of the heat input from the chickens and the high water activity of the droppings, but occasionally, if the shed is inadequately ventilated, the water builds up and the litter becomes uncomfortably saturated.

Heat also passes through to the concrete and the underlying soil. Since the shed is virtually in thermal equilibrium at all times, a simple calculation revealed that roughly

1W per chicken is conducted through to the ground. Similarly, in climatic conditions typical of Auckland, roughly 1W per chicken is conducted through the shed's wall and roof. This leaves about 8 W per bird, or a total of about 240 kW of heat to be removed by ventilation.

The group made a sample calculation based on an air speed of 1 m/s provided by the ventilation fans. With ambient (outside) air conditions of 20 °C and a relative humidity of 60%, the expelled air was calculated to be at 20 °C with a RH of 70%. This provided a water uptake of 0.1 kg/s and a heat gain of 235 kW, which matched the sample data closely (see above). This was very encouraging, as it showed that a simulation of the thermodynamics and psychrometry of the shed environment produced feasible results.

The rate of food and water intake was also investigated using Tegel's data. We found that the chicken's intake was proportional to their surface area, assuming that they were spherical – remarkably, this latter assumption appears to be fairly good! The rate of growth appeared to be quadratic with age. A model for the heat production and water respiration rate of a typical chicken was developed, based on physical principles. Perusal of actual data on chicken weight vs feed and water intake led to some initial simple models for growth rate as a function of mass. These will be refined for the final write-up.

Industry representative John Foulds said that the models developed by the team confirmed Tegel's thoughts about important parameters, but he was impressed by the way the group attacked the problem and how the members insisted on ensuring that everything was accounted for.

# FORCASTING WIND FARM GENERATION

## Transpower

### Industry Contacts:

David Rohan  
Conrad Edwards

### Moderators:

Dr Tasos Tsoularis, Massey University, Albany  
Dr Bill Whiten, University of Queensland

Wind power is an important potential source of large amounts of renewable energy. Unlike conventional energy generation, the amount of energy produced is not able to be specified in advance as it depends on the wind velocity which can vary significantly over short periods of time. This makes balancing supply and demand more difficult. An ability to predict the possible range of wind velocity for the next 5 minutes to the next 24 hours would assist in the scheduling and purchasing of power. Power supply is balanced with demand on a five- minute basis and bids for purchase of power become firm two hours ahead. The uncertainty due to large amounts of wind power on a power network is expected to make both balancing power supply and demand, and realistic bidding for supply much more difficult.

This MISG project looked at the prediction of wind velocity and associated power generation to determine how well these can be predicted. The Tararua wind farm supplied scaled wind and power data from their site, and meteorological data and predictions had been obtained for the three nearest cities.

An initial graphical analysis was undertaken to get an understanding of the available data. This determined the periods for which data was available and how missing data had been coded. Significant differences were found in the data from the two sources. The data was then examined using several techniques.

The simple persistence method, which assumes a constant wind velocity is used as the basis to compare other methods. It was found this method gives good results for prediction of short time periods but loses accuracy as the prediction time increases.

Linear regression was used to develop simple prediction equations. These clearly showed a progressive loss of accuracy as the prediction time increased beyond about three hours.

The Kalman filter is a linear technique that progressively updates a prediction equation according to the error in the last predictions. An initial test of a Kalman filter was undertaken. This technique is considered to be promising as it has the potential to adapt to changing conditions.

Several neural network techniques were investigated. Neural networks provide a black box method that can conveniently develop predictions of complex responses behaviour. However developing (training) a neural network requires much more computation than the linear methods above. The initial neural networks were able to some improvement over the persistence method. Specifically, a study by Zeke Chan of AUT indicated that ANNs reduced the average forecast error by 13% in a 4-hour forecast and 16% in a 12-hour forecast. Ray Hoare's Multi-Layer Perceptron (MLP) study showed that standard neural network software could be used to produce useful predictions.

An extension of networks due to Timothy Hong was also tested. This divided the dependent variables into regions using multiple fuzzy sets and trained a network for each region. Predictions were made by combining the neural network predictions according to the fuzzy set memberships. Again improvements over the persistence method were demonstrated.

All the methods investigated potentially improve on the simple persistence method, however none gave large improvements. According to EU experience a combination of prediction approaches seems to be most promising. It is firmly believed that past data from one site cannot provide accurate longer-term prediction of wind, and data from other sites will be needed to improve such forecasts. It is recommended that an individual model for each wind farm site needs to be developed. Meteorology data is collected on a world wide basis and together with forecasting techniques provides the information from other sites processed to provide local weather forecasts and thus has the potential to improve on the longer-term forecasts.

# Earthquake damage in underground roadways

## Solid Energy NZ Ltd

### Industry Contacts:

Richard Mould  
Stephen Bell

### Moderators:

Professor Tim Marchant, University of Wollongong  
Dr Graham Weir, Industrial Research Ltd

Solid Energy operates the two coal mines, Terrace and Spring Creek, on the West Coast of the South Island. They asked the MISG group to quantify the damage likely to occur in their mines due to a magnitude eight earthquake on the Alpine fault. The mine workings are typically 200-400m deep and they are about 40 km distant from the Alpine fault.

Firstly, the group investigated the wavelength and type of waves likely to be incident upon the mine workings. Using published data, it was determined that the most energetic wavelengths in the earthquake response spectrum were 200 to 500 metres in length. Also, the group calculated that Rayleigh, or surface waves, decay to one per cent of their surface energies at a distance approximately 100 metres below the surface. A survey paper in the literature indicated that of 132 cases of earthquakes at mine sites moderate or heavy damage rarely occurred below 100 metres. This result fits nicely with the decay scale of the Rayleigh waves calculated by the MISG group. Hence Rayleigh waves have no impact and the mine workings are subject to long S and P body-waves only.

Secondly, an empirical relationship from the literature was used to determine the peak acceleration of the waves at the mine workings, in terms of the earthquake magnitude and the distance from the fault. For the Terrace and Spring Creeks mines peak accelerations were estimated to be 0.2 g, where g is the acceleration due to gravity. The literature suggests severe damage occurs for accelerations greater than 0.5 g; hence it is likely that the Terrace and Spring Creeks mines would be subject to light damage only.

Thirdly, the interaction of S body-waves and the mine roadway was considered numerically using a finite-element package. The roadways can be 500 metres long, which is comparable to the wavelength so the possibility of resonant interactions was investigated. The numerical results indicated that some slight amplification of the strain did occur, due to the presence of the roadway, but resonant amplification did not occur.

Finally, consideration of the energy released by a magnitude 8 earthquake on the Alpine Fault suggested that this may occur either as one event (as is usually assumed), or as a series of discrete events, analogous to domino collapse. The latter

scenario could reduce the calculated damage, especially in the near field.

In summary, serious damage to the Terrace and Spring Creek mine workings, which are 200 to 400m deep, is unlikely to occur as the result of a magnitude eight earthquake. Rayleigh waves are likely to damage the mine portal however, and damage to the surface portal seems a good area for future investigation.

# Modelling the spread of wilding conifers

## Environment Canterbury

### Industry representatives:

Philip Grove  
Yvonne Buckley

### Problem moderators:

Associate Professor Mick Roberts, Massey University, Albany  
Heather North, Landcare Research

### Problem overview

As a Regional Council, Environment Canterbury is responsible for many aspects of environmental protection and sustainability. A major weed problem has been recognised in the form of wilding conifers spreading from existing plantations and shelterbelts. These threaten native vegetation and important wildlife habitat, as well as impacting on pastoral farmland and the visual and recreational values of Canterbury landscapes. A recent survey by Environment Canterbury mapped over 60 thousand hectares of conifers, including both plantations and wildings. It is estimated that seed dispersal from these threatens a further one million hectares of land. Under the Regional Pest Management Strategy 2002, Environment Canterbury has some funding for wilding conifer control, but this is limited relative to the scale of the problem. Thus, the control budget needs to be strategically allocated to achieve the best possible environmental outcomes for the dollars spent. The problem posed to the MISG centred on modelling wilding conifer spread, with the overall aim of prioritising sites for control operations.

### Progress at the MISG

The study group focused efforts on three interrelated aspects. The first was to develop an understanding of the topographic and climatic drivers of short and long distance seed dispersal, in order to identify which existing conifer sites posed the greatest risk to the surrounding land. Fringe spread occurs when seed is released from a tree and is carried some horizontal distance by the wind, until it falls to the ground under the force of gravity. However, *Pinus* seeds will typically only spread about 100m in this case. To achieve the 8–10km dispersal distances observed in invasive field situations, the group showed that the effective release height of the seeds had to be much greater than tree height. The literature suggests that thermal uplift may be a major mechanism for this height gain. However, this seemed unlikely at the site for which we had data at MISG – the Mount Barker site, near Lake Coleridge in the Canterbury high country. Strong North West winds are the norm when the temperature is high, so thermal uplift is probably rare. The group thought that Mount Barker itself was functioning as a launching ramp for seeds released from mature trees upwind of the hill. They could be carried by the wind up and over to be effectively released at a height 200m above the surrounding land. Even though the great majority of seeds would be dropped in the area of low pressure in the lee of the

hill, some would continue on the airflow, and could travel the 8km distance on winds over 100kph. Why then are these long distance spread events so rare – for example, only one to three major events at Mount Barker in 100 years – if topographic uplift can occur so easily? Some initial calculations showed that wind at slower speeds would tend to go around, rather than over, the hill, so seeds may not obtain the required release height. To assess risk of spread from a site, the group recommends that topographic and climatic factors be analysed to estimate both the maximum possible spread distance and the likely frequency of long distance spread events.

The second area of effort at the MISG was modelling the dispersal of seed once it has been released. There are several distributions suggested in the literature for modelling the density of seed rain and the distances travelled. These models describe the high seed densities that fall out of the airstream close to the source, and the long tail of seed that travels the greater distances. It is the outermost edge of the tail that determines the invasion speed of the conifers. Thus it is important to correctly model the distance and seed rain density for those seeds that travel the farthest, in order to correctly predict invasion speed. The group worked with a set of data from Mount Barker, which gave distance-from-source for each tree in the down-wind tail of wilding conifers. Several models were fitted to the data, and the group found that a model first proposed by Okubo and Levin described the tail well. The assumptions and approximations made in this model must now be checked before it can be recommended as appropriate for studying wilding conifer spread.

Finally, the group assessed the effectiveness of a range of management options for controlling conifer spread. This analysis was carried out by first developing a pragmatic model of invasion as a series of discrete steps occurring in a down-wind direction. The group defined a set of possible transitions from young to old and scattered to dense trees, and assigned probabilities to these. This enabled a computer simulation of year-on-year invasion to be developed. Knowing that there was not sufficient budget to remove all trees, the task was to find the management strategy that provided the best control of invaded area. A range of possible management strategies were defined, such as always targeting the largest patches of trees for removal, or always targeting the oldest. These were also simulated, with the chosen management strategy of removing particular patches of trees each year, and the remainder allowed to continue growing and spreading. Some strategies, such as targeting the oldest trees, did not provide good control. Targeting the youngest (which would equate to the outermost edges of the tail of wilding conifers) worked much better. The number of trees cut each year could also be varied in this simulation. The next steps would be to bring in a more detailed model of spread (relatively simple assumptions were used during this initial development) along with dollar-value costs for each of the control strategies. This simulation has the potential to provide a valuable framework for predicting and visualising the effect of various control strategies.

## **Conclusions**

The group at the MISG has increased understanding of topographic and climatic drivers of wilding conifer dispersal distances and event frequency. A more realistic model of seed rain density has been tested, supporting more accurate prediction of invasion speed. These developments are important for site risk assessment. A framework has been developed for modelling conifer invasion in space and time, and for predicting the level of control expected from various management strategies.

# Optimal sorting of product into fixed weight packaging

## Compac Sorting Equipment Ltd

### Industry Contacts:

Mr Nigel Beach  
Mr Gavin Reeve

### Moderators:

Dr Clive Marsh, Canesis Network Ltd  
Dr Philip Kilby, CSIRO, Canberra

Compac Sorting Equipment make very nifty machines for sorting fruit by weight, diameter, colour, density, blemish or even shape. Compac sought solutions to two closely related problems: the boxing problem and the bagging problem. The boxing problem requires graded fruit to be assigned to outlets where boxes are filled with a specified number of fruit to a minimum weight (and a specified tolerance for underweights). The aim is to maximise the number of boxes packed. The decision must be made after all information is known, but before the fruit passes the first outlet – a few seconds total. Further, information about fruit already packed in a given box is incomplete (we don't know exactly which fruit ended up in a box). The bagging problem requires bags to be filled to a minimum weight – no tolerance for underweights, and no constraints on the number of fruit per bag. In this case complete information is available on fruit already assigned to a bag. Again the aim is to maximise the number of bags packed.

The MISG team were able to provide a 'close to optimal' solution to the boxing problem for the simplest information scenario where an irrevocable decision is made for each fruit in turn, and no memory of previous assignments is kept. This information scenario is the least demanding of real time measurement for Compac and is also the simplest to analyse and optimise. Basically, it is an attempt to improve on an old idea already implemented by Compac – static cut points. Cut points define the category (and hence the outlet) a fruit will be assigned to. Potentially, the value of the fruit may be different for different categories. The cut point optimisation problem was formulated with the objective of revenue maximisation. The determination of globally optimal solutions is non trivial but some approximations and assumptions yielded a tractable solution method. We are confident that the solutions given by this approach won't be substantially worse than those from any other approach; hence we claim a 'close to optimal' result. Compac believe that there is value in adapting the choice of the cut points to the distribution of the incoming fruit. The solution method proposed does depend on this distribution and it can be re-run if the distribution changes significantly. We were able to provide software to solve the cut-point optimisation and a simulation of the resulting decision system. The simulation results compared favourably with Compac's current approach.

The team investigated a range of approaches to realise the potential benefits of richer information scenarios. Each method assumed some knowledge of previous assignments to an outlet. They also exploited the fact that the system has some time

between when all information about the fruit is known, and when it must commit to a decision. Some tens of fruit can therefore be considered together. These approaches use the fact that weight allowances of fruit categories overlap. An 18kg, 100-fruit box has an average weight of 180g per fruit. But if we aim for a target weight of 180g, the variance in individual weights means we will often get underweight boxes. The proposed methods choose target weights sufficiently above the minimum target (180g in this case) using the variance for weight of fruit already sent to the outlet so that it is unlikely a box will be underweight. The trick then is to minimise the variance. One of the algorithms looked for good combinations of pairs of fruit (even pairs of pears).

It was not possible to fully quantify the improvements that use of this extra information would yield but a range of algorithms were developed and can be investigated further.

We looked at the bagging problem in two ways. The first concentrated on the physical aspects of the problem. The physical constraints have a very large impact on the problem. For example, fruit from a near-side lane arrives faster than far-side lanes, so if we are just finishing off a bag, near-side fruit is much preferred. Restrictions in the number of fruit that can be allowed to cycle must also be observed. A Matlab model was created that simulates some of these physical restrictions, and has been made available to Compac. An algorithm for choosing the fruit based on “preference zones” was developed, where bags close to finishing were allowed to choose fruit from their “most-preferred” zone before other bags got to pick over the remaining fruit. Unfortunately, we didn’t have enough time to run it through the simulation.

A second approach concentrated on trying to find the best three or four fruit to finish off a bag. A method was presented that is feasible for real time implementation both in terms of the information (measurements) and computation required. Simulations suggested that this algorithm significantly out-performs the current, simplistic approach of first-in-first-out. It currently assumes pre-graded fruit but could be made more generic – allowing fruit that can be classed into more than one grade.