Proceedings of the 2020 Mathematics in Industry Study Group

Natalie Thamwattana[1] Michael H. Meylan[2]
A.J. Roberts[3]

2022-12-06

Abstract

This special Section of the ANZIAM Journal (Electronic Supplement) contains the refereed papers from the 2020 Mathematics in Industry Study Group (MISG2020) held at the University of Newcastle from 28 January to 1 February 2020. This report provides the equation-free outcomes.

Contents

doi:10.21914/anziamj.v62.17406, © Austral. Mathematical Soc. 2022. Published 2022-12-06, as part of the Proceedings of the 2020 Mathematics and Statistics in Industry Study Group. issn 1445-8810. (Print two pages per sheet of paper.) Copies of this article must not be made otherwise available on the internet; instead link directly to the doi for this article.

1 Preface to the proceedings

This special Section of the ANZIAM Journal (Electronic Supplement) contains the refereed papers from the 2020 Mathematics in Industry Study Group (MISG2020) held at the University of Newcastle from 28 January to 1 February 2020. The MISG is a special interest meeting of ANZIAM, the Australia and New Zealand Industrial and Applied Mathematics (ANZIAM) division of the Australian Mathematics Society. The MISG meetings take place annually and provide a forum where projects proposed by industry can be worked on intensively, by high profile scientists in the fields of Applied Mathematics, Statistics and Engineering, from Australia, New Zealand and the world beyond, along with representatives from the industries proposing the projects.

The writing of these papers was coordinated by the project moderators in consultation with the coauthors and company representatives. The manuscripts were submitted to the editors, Associate Professor Mike Meylan, Professor Ngamta Thamwattana and Professor Tony Roberts, and were subsequently refereed by two expert referees. On the advice of the referees, manuscripts were accepted for publication, subject to the recommended revisions, and formally approved by the editorial committee.

At MISG2020, four projects were presented from diverse industries, with 78 delegates participating.

1.1 Industry Partners

We gratefully acknowledge the support of our industry partners:

- Lovells Springs;
- Safearth;
- Concrush;
- Hyper Q Aerospace.

1.2 Acknowledgements

In addition to our industry partners, we gratefully acknowledge support from the following organisations:

- ANZIAM;
- Office of the NSW Chief Scientist & Engineer, Department of Industry, NSW Government;

- Priority Research Centre: Computer Assisted Research Mathematics and its Applications, The University of Newcastle;
- Faculty of Science, The University of Newcastle.

We are also grateful to Professor Ryan Loxton from the Centre for Optimisation and Decision Science, Curtin University, for giving a public lecture on power of optimisation research in mining, energy, and agriculture industries, as part of the MISG's outreach event and acknowledge the support from the Hunter Branch of the Royal Society of NSW in promoting the public lecture. We are also grateful to Professor Mark McGuinness (Victoria University of Wellington), Professor Troy Farrell (Queensland University of Technology), Associate Professor Amie Albrecht (University of South Australia) and Dr Neville Fowkes (University of Western Australia) for their helpful advice and comments in organising the MISG2020.

1.3 MISG2020 Organising Committee

- Professor Ngamta Thamwattana (Co-Director)
- Associate Professor Mike Meylan (Co-Director)
- Mrs Juliane Turner (Administrative Support)
- Dr David Allingham (Technical Support)

2 Safearth non-technical summary

Moderators: Glen Livingston, Geoffrey Pritchard, Riya Aggarwal, Ian Griffths

It is well known that under strong enough voltages/currents, the human heart will stop beating, an effect known as ventricular fibrillation. This is the principal danger of receiving an electric shock. Currently, the safety standards in Australia specify a safe voltage/current, below which the effect is considered not to be life-threatening. In practice, the individual human response to voltages/currents varies greatly, and it would be better to develop a more sophisticated model in which the likelihood of injury was modelled. Due to the desire for large safety margins, the safe voltage must be one for which the probability of ventricular fibrillation is minimal.

Recently, Safearth developed a safety model in which they propose a calculation framework for quantifying, given a shock hazard, the probability that an individual randomly chosen from the population will experience ventricular fibrillation. They wanted to know whether this model is the best way to approach the problem or whether other methods should be considered.

In the situation modelled, a given 50 Hz AC voltage is applied across the body (via a specified current path, such as hand-to-hand or hand-to-foot) in series with another impedance (representing shoes, gloves, soil, etc.) for a given length of time. This defines the shock hazard. The fate of the individual then depends on two key characteristics: the body's impedance (which determines how much current will flow in the circuit) and the heart's current tolerance (the minimum current required to trigger ventricular fibrillation).

2.1 Body impedance

The impedance of the human body is a decreasing function of the applied voltage. This is largely due to the behaviour of the skin, which presents a relatively large impedance to low voltages, but breaks down for voltages on the order of a few hundred volts. The distribution of body impedances across the human population will thus also be voltage-dependent; the available data consists of the 5th, 50th, and 95th percentiles of this distribution for a set of about thirteen specified voltages ranging from 25 V to 1000 V.

Safearth's original model posits that each individual's body impedance coincides with the same population quantile at all voltages. (So, for example, an individual whose body impedance at 25 V happens to match the population median will have a body impedance at 1000 V matching the population median for that statistic also.) In effect, this gives a one-parameter family of body impedance functions: any individual's body impedance is completely specified, as a function of voltage, by the population quantile with which the individual coincides. Safearth assumed the distribution of body impedance to be lognormal for each applied voltage.

A slightly more complex model developed at MISG considers the skin and the subdermal tissue separately, as two impedances in series. Only the skin impedance is voltage-dependent. In effect, this gives a two-parameter family of body impedance functions. The two impedances are assumed to be independent and both lognormally distributed. The impedance curves of two different individuals may cross, something not possible in the Safearth model. One merit of this model is that the impedance function of any individual must be decreasing with voltage; this is not guaranteed (and in fact does not hold) for the Safearth model.

2.2 Current tolerance

An individual's current tolerance depends on the exposure duration; prolonged electric shocks are more hazardous than brief ones. Most previous work, including the basic Safearth model, has assumed the distribution of current tolerance to be lognormal for each duration. The MISG group largely accepted this assumption, and focused on finding functional forms to describe the dependence on duration. A logistic function was considered, but the fit was found to be not particularly good. An alternative approach fitted a cubic spline; this was found to work well with only a modest number (thirteen) of degrees of freedom.

2.3 The shock circuit with additional series impedance

An individual's current tolerance and body impedance function are sufficient to determine the probability that ventricular fibrillation occurs when the body is subjected to a given voltage for a given length of time. However, the voltage to which the body is exposed (touch voltage) may not be the full source voltage of the circuit, due to additional series impedances arising from protective clothing, shoes, or the soil.

Since the body's impedance is (non-linearly) voltage-dependent, finding the touch voltage requires solving a nonlinear equation. As Safearth had already discovered, the solution can in most cases be found by fixed-point iteration. However, the MISG group found some cases in which this method does not converge to a solution. As an alternative, a more straightforward bisection search can be used.

3 Hyper Q Aerospace: Aerodynamics and Control of Next Generation Electric Rotorcraft

Moderators: Peter Batten, Jim Azar, Graeme Hocking, Winston Sweatman, Benjamin Maldon.

The company "Hyper Q Aerospace" is developing an aircraft with two coaxial rotors that will use electric axial flux motors to rotate the rotorhead and attached rotor blades. There will be no need for gearboxes, transmissions and drive shafts such as in conventional helicopter design. The major factors to be considered were the vibrations in the system, the noise generated by the new design and the aerodynamical response. These questions are considered and we summarise the findings of each in turn.

Vibrations are an inevitable occurrence in systems of this type. In general, these vibrations are harmless to the integrity of the craft unless there is some feedback between components that leads to a resonance and consequent growing oscillations. In that case, serious damage may result. Possible sources of vibration are the blades as the velocity changes, the movement of the entire rotor, and the interaction between the rotor disks and the blades as they cross.

The frequency of all of these sources of vibration can be readily determined with the exception of the rotor blades. The vibration characteristics of the rotor were calculated by considering its behaviour to be similar to that for a general beam using the Euler-Bernoulli equation. This equation models the oscillations in a beam of small thickness by simplifying the equations of solid mechanics. So long as the oscillations are relatively small it provides an accurate representation of the behaviour of the beam given the basic material properties, such as Young's Modulus E and moment of inertia I, for a beam of length l and mass m. Along beam oscillations and across beam oscillations were considered and the resonant frequencies and harmonics were computed.

The nice thing about this approach is that the solutions can be computed quickly and exactly given the characteristics of the blade.

This calculation provides the major frequencies within the rotor that should be avoided in other components within the rotorcraft, such as the engine vibration of rotor cycling or rotor crossing, since if these are kept at the same value of frequency the energy within these may feed into the rotor oscillations causing resonance and ultimately failure.

3.1 Noise Footprint

The noise generated by a helicopter was considered via a literature search and the findings indicate that most of the noise comes from only a few sources. One is the rotor itself, another is the interaction of the rotor with the tail-rotor (on single main rotor craft), engine noise and other mechanical noise. The largest components of the noise are in the lower frequency sounds generated by the rotor itself.

In the design of the Hyper Q Aerospace helicopter, much of the higher frequency noise will not exist. This is due to the nature of the electric motor and direct drive on the blade system. The lower frequency noise will come from the rotation of the rotor and the crossing of the blades, but it is difficult to determine the volume of this noise without further work. The main component is likely to be due to the bump of blades crossing and so having aligned blades on the rotor will almost certainly amplify the volume.

3.2 Lift

One of the most important components of a helicopter design is the lift generated by the rotors. In the design of Hyper Q Aerospace, the rotation speed is envisaged to be between 250–1500 rpm. The group considered the lift on the blades and the effect of blade crossing and the resulting “bump”.

Flow past a single airfoil was considered to determine the region-of-influence in the surrounding air. This can be observed as the region over which the air flow lines (streamlines) are disturbed in steady flight. The work of aerodynamics pioneers including application of the Joukowski transformation, thin airfoil theory and Rankine bodies is considered. These classical results produce very accurate computations for lift on an airfoil of any reasonably small aspect ratio. The results provide a good approximation to the local flow without the undue complication of a full computational fluid dynamics simulation. Each blade of the helicopter produces an average lift L, and so the total lift on a single rotor helicopter can be accurately ascertained by summing these lift calculations over all lifting surfaces. Close to the ground, the blades generate extra lift due to the effect of the pressure increase obtained as the downdraft hits the ground. This is known as "ground effect" and only occurs when the blades are very close to the ground.

The traditional method of computing steady flows is to consider the flow in the frame of reference of the airfoil. However, for a helicopter with coaxial rotors moving in opposite directions this reference frame is not applicable because the problem is changing with time (unsteady).

Another related flow is that past the wings of a bi-plane. Although the wings of a biplane are travelling in the same direction, unlike the rotor blades considered here, the bi-plane does provide information about the passage of air between parallel aerofoils. Bi-planes are generally accepted to produce about 1.4 times the lift of a single wing aircraft. So although there is extra lift in comparison, the bi-plane does not provide as much lift as that due to two separate single wings.

The full aerodynamic problem is very difficult but the group felt that after considering what had been learned during the week some general conclusions could be drawn. When the blades pass over and under one another the two interact in a manner that can be thought of as a "ground effect". The ground effect on the upper blade due to the lower blade increases its lift due to an increase in pressure beneath. However, this pressure increase above the lower blade reduces the lift on that surface. To some extent these effects will cancel during a hover, while when moving forward the blade that is moving toward the direction of travel will experience the greatest change. It seems likely that with the extra impact of induced drag between the rotor blades, there will be some loss of the total lift generated by the two lifting surfaces. If there are the same number of evenly spaced blades on each rotor, then all of the blade crossings coincide and the effect will occur simultaneously. If the upper and lower rotor blade numbers are coprime, then the blade crossings will all occur separately, perhaps producing a smoother lift profile during a cycle.

3.3 Final Comments

The study group considered several factors in the design of the helicopter. It was possible to come up with a formula for the vibration modes of the rotor blades that should be avoided. Analysis of the noise footprint will require further work, but it is likely that this will depend on the configuration of the rotor blades. Lift was determined to be a function of rotor speed, separation and blade interactions. Speed and area increase lift, but blade crossings may decrease it. It would seem likely that greater vertical separation of the rotors will allow them to interact less leading to slightly greater lift on average for each rotor. It is likely that there is an optimal combination of all of these factors.

4 Equation-free summary of dust emissions from concrete recycling

Moderators: Brendan J. Florio, Fillipe Georgiou, Melanie E. Roberts, Kevin Thompson.

Concrush is a concrete recycling operation located in Teralba in the Hunter region of New South Wales. Essentially, the plant takes in concrete waste and crushes it. The crushed product is then on-sold as roadbase, drainage aggregate, and perhaps in the future, aggregates to produce recycled concrete. The targeted size of the crushed product is determined by its intended use.

For example, drainage aggregate measures up to 20 mm while roadbase requires a spectrum of particle sizes as a mix of course and fine aggregate allows the desired compactability properties. In addition to concrete, Concrush also process greenwaste to produce a mulch product. As a by-product of the concrete processing, very fine "fugitive" dust is produced, which can become airborne and contribute to poor air quality.

The fine particles of interest fall into two categories: PM10 and PM2.5. These correspond to particles of diameter $< 10\,\mu\text{m}$ and $< 2.5\,\mu\text{m}$ respectively. Particles that span less than 10 microns are able to invade the upper airways and can cause irritation while particles that span less than 2.5 microns can penetrate the lower airways and even enter the bloodstream. Irritation of the lower airways by PM2.5 particles may be problematic, especially for asthmatics, and long term exposure is not recommended.

Air quality in the Hunter region is historically a problem, and the NSW Environmental Protection Authority (EPA) have set limits on the amount of fine particles generated at industrial sites such as Concrush. As a small business of the Hunter community, Concrush is driven to reduce the emission of fugitive dust for both regulatory reasons and a sense of community responsibility. Currently, Concrush maintain four dust gauges on the perimeter of the site. These are analysed monthly for fine dust content to give an insight to their dust emissions.

The guidelines imposed by the EPA are the same as those that are imposed on the coal industry. Are these regulations appropriate for a concrete recycling plant? If they are too strict, then they may unduly effect the productivity of the site. If they are too loose, then significant dust generation may occur despite regulatory approval. Consequently, the models and assumptions of the EPA may be inappropriate, so Concrush are seeking external guidance at the Mathematics in Industry Study Group. Concrush are seeking a greater insight into the production and dispersion of dust to drive suppression strategies. Based on the overall context, we distilled the problem into three main questions:

- What are the critical processes that produce airborne dust and how much do they produce? The traditional culprits are
  - The crushing and grinding processes,
  - Erosion of fine dust from aggregate stockpiles and unsealed roads,
  - Wheel-lifted dust from the transit of large vehicles over unsealed roads.
- Of the dust collected in the perimeter gauges, what proportion of this dust is due to Concrush operations? What proportion is due to neighbouring industries? Concrush is surrounded by
  - A freight line where coal is transported,
  - A high traffic main road,
  - A former lead mine.

  These may have a significant contribution to the local air quality.
- Once the dust is produced, where does it go? How far does it travel? This would also help determine the sources of dust beyond Concrush.

We chose to address this problem with three different approaches: data analysis, a literature review and mathematical modelling. In the following sections we summarise our activities in each of these three approaches. For data analysis, we use data from the dust gauges, a handheld indoor air quality unit and nearby measurements from the Department of Planning, Industry & Environment. Dust emission is a problem in many different industries and has been studied extensively. A review of available literature would assist us in obtaining parameter values for our mathematical models as well as providing several rules-of-thumb to suppress dust emissions. Finally, mathematical models can provide quantitative and qualitative insight into the transport and generation of dust.

4.1 Data analysis

The group considered a variety of onsite and offsite data sources. It is difficult to draw any meaningful data from the onsite dust gauges as they are measured monthly. At this point, we look further afield for reliable data.

The Department of Planning, Industry and Environment have published offsite data from three nearby stations, including one within 10 km of the Concrush site. One set of data is interesting: The five year time of day average Parts per Million (PPM) value for PM10 particles. The resulting data is bimodal with distinct peaks at about 8am and 6:30pm, which we believe corresponds to peak hour traffic times. This could be a contribution from wheel-lifted dust, or diesel emissions of which PM10 particles form a minor population peak.

PM10 concentration was continuously monitored onsite near the weighbridge using a handheld monitor. Data was collected every five minutes over a period of fifteen days. Preliminary analysis seems to suggest that there is significant causality between truck movement times and PM10 spikes. Concentration spikes occur significantly 2–5 minutes following truck activity. Causality was determined by the Granger test.

The main problem with the available data is that most of it is given in monthly aggregates. While this may be satisfactory for regulatory reasons, it is almost useless for us to analyse the emission and transportation of dust. Given that continuous monitoring is available, our group suggests that is is put to use longer term on site. It is recommended that multiple sensors be deployed to provide spatial data. This would, for example, go a long way toward isolating the emissions of the Concrush site compared to the background air quality. Measurements taken during site downtime will be almost entirely background dust, with only the passive emissions from stockpiles complicating the matter.

4.2 Literature review

Many aspects of the dust emission problem are unknown. In particular, what is the raw mass of dust produced for a given time in any of the recycling processes or other activities conducted on site? The EPA guidelines are formed based on coal dust, however their applicability to concrete dust warrants further investigation.

Our preliminary search has uncovered a few qualitative maxims regarding dust production. While the actual crushing of concrete is the only process that actually produces fine particles, the emission of fugitive dust during this process is quite low as it is well-shielded. Similarly, the passive emission of dust from stockpiles is also low except during high wind. Finally, moving material around the site and dumping in particular, causes a moderate emission of fugitive dust.

There are two main processes undertaken by Concrush that contribute to large emissions of fugitive dust. The first is truck movements, particularly over unpaved roads. The floor of the site is covered by accumulated fine particles which have conveniently settled. Truck movements disturb these fines and can send them into the air as fugitive dust. The second high-emission process is screening. Crushed aggregate is passed through a series of agitated sieves with progressively finer grading, which acts to sort the aggregates. As this is a mechanically energetic process, so fine particles may become airborne and escape as fugitive dust.

Wind is a significant factor in both the emission and transport of dust. In fact, Concrush cease all dust producing operations if the local wind speed exceeds 18 km/hr. Background wind draws fine particles out of stockpiles, so the emission of dust increases with increasing wind speed, usually above some threshold value. Simple strategies to reduce wind-generated emissions include the following.

- Building barriers to reduce wind speed. These can range from simple windbreaking cloth, to a fully enclosed building.
- Wet stockpiles during high winds. Concrush already has the facility to wet the stockpiles and they do this for product quality reasons. If water resources allow, continuing this process during high winds would lower dust emissions.
- Arrange stockpiles such that piles of coarse aggregate act as a windbreak to protect piles of fine aggregate. Such arrangements should take seasonal dominant winds into account.

Any method to reduce the wind speed by arranging windbreaks would go a long way to reduce dust emission, due to its relationship with both emission and transport. At the moment, the Concrush site is completely open, so this is an easy area for improvement.

4.3 Mathematical models

To describe the transport of fugitive dust particles, the group turned to mathematical models. We considered a wide range of approaches, including an advection-diffusion model at the simple end. To capture more complex behaviour, models from other applications were adapted. Namely, we considered adapted versions of a saltation model and a bushfire plume model.

The advection-diffusion model contained two basic phenomena: the advection of dust particles by background wind, and the dispersal of dust particles. Note that the "diffusion" we consider is not molecular diffusion, but the general mixing behaviour caused by turbulent eddies, for example. For a first pass, we assume that a certain amount of dust is released in a narrow vertical column, and then assume that the wind and dispersal only act horizontally. This allows us to obtain an exact 2D solution by taking advantage of some properties within our model. This solution describes the dust concentration at any given point at any given time in the surrounding area.

We also took into account the settling and deposition of dust particles onto the ground surface. The vertical settling of small dust particles can be set by the Stokes velocity, which describes a linear loss of mass over time. The model can accommodate classes of different particle size, and as a result, we can see that larger particles fall swiftly, while small particles fall slowly and may travel far with the help of wind. Computer-based simulations were also implemented and verified against the exact results. This allows us to extend the simple model with additional physical phenomena and obtain a numerical solution. Non-point sources, wind shear and particle interference would all require numerical treatment

We considered a saltation model to describe the motion of particles as they "skip" along the ground. Such models are traditionally used to described the erosion and evolution of dunes and other sand patterns on desert landscapes. This model gives us insight into the behaviour of different sized particles, how they interact with the ground and how they are influenced by wind and turbulence.

We also investigated the utility of available fire plume models. These models are excellent at predicting particulate levels from aggregates sources over a large geographic area. Ideally we would characterise the geography surrounding Concrush to determine the source of both background and emitted dust levels. Realistically, we nee more data.

4.4 Conclusion

The study group was an overall success. The group tackled the dust problem on three fronts to answer the questions of how the dust is produced, where it travels, and who is responsible for it. Many results require a quantitative description and are not appropriate for this equation-free summary.

On the qualitative front, based on literature review, the group identified that wind is a major factor in dust transportation and emission. Any sort of windbreak should go a long way in reducing the flux of dust from the site into surrounding areas, and that is our primary recommendation.

Acknowledgements The moderators thank the industry representative, Kevin Thompson, for being available for the entire week. We also acknowledge the hard work and contributions of the group members: Susam Boral, J. Divahar, Olivier Huet, John Knight, Cameron Lord, Ognjen Orozovic, Björn Rüffer, Tarun Srivastava, Matthew Tam, Thien Tran-Duc, Kyle Stevens, Dimetre Triadis and Jennie Trumpeta.

5 Lovells Springs

Moderators: Simon Crane, Mark McGuinness, Barry Cox, Balaje Kalyanaraman.

Lovells Springs are planning to construct a new furnace to heat rods to workable temperature prior to coiling into springs. The furnace is heated by a small number of natural gas jets which are aligned horizontally in the upper section of the furnace. The jets are coupled with thermostats and the flow of gas and air is independently modulated to maintain a constant furnace temperature of around 1000°C. The furnace interior is approximately 1.5 m high, 2.2 m wide and of a length prescribed by the rod lengths used in the furnace and also whether there will be coiling happening at one or both ends of the furnace. Typical lengths are 5–8 m. Figure 1 shows photographs of a furnace currently in use at Lovell's.

The articulation of rods through the furnace is achieved with a system of notched blocks that are coupled with an hydraulic walking mechanism. The hydraulic system can complete one cycle approximately every 15 seconds. The sizing of the scallops is determined by the rod diameter and it is anticipated that the proposed furnace would be used for the manufacture of rods into springs with a rod diameter in the range 14 mm to 24 mm.

A large component of the operating expense is the cost of the gas firing the furnace. Therefore any improvements to processes that reduce gas usage would have the dual benefit of making the business more economical as well as reducing greenhouse gas emissions. The stepping of rods through the furnace is aimed at minimising carbon loss from the steel and to have the rod exit the furnace at an appropriate working temperature (around 950°C) and at a manageable rate (around one rod per minute).

The time required to heat a rod is primarily a function of the diameter of the rod but also depends on the initial temperature of the rod when it enters the furnace, the temperature of the furnace, the rate at which heat

Figure 1: A furnace viewed from the vestibule side, and a close-up view of hot rods through a side hatch of the main furnace region.

energy is input to the oven, and the emissivity of the rod. This final point is particularly salient because many rods are pre-peeled prior to heating which makes them shiny and decreases the emissivity, slowing down the transfer of heat into the steel rods.

Lovell's Coils seek guidelines for the design of the vestibule at the furnace entrance. The vestibule is envisioned to be the main outlet for the flow of hot exhaust gas exiting the furnace over the rods which are approaching the furnace. The purpose of the vestibule is to pre-heat the rods in contact with the exhaust gas. A reasonable target may be to raise the temperature of the rods to 300°C prior to loading in the furnace proper. In particular, Lovell's Coils would like to see a heat transfer model that predicts the effects of varying the width and height of the vestibule. They are also interested in the relative importance of radiative versus convective heat transfer mechanisms.

5.1 Modelling Overview

The group developed a variety of models for the transfer of heat into the rods in the vestibule. There was a focus on the relative importance of radiative heat transfer to the rods in the vestibule, and convection/conduction mechanisms for heat transfer to the rods from the hot gas flowing past them.

A system of two coupled convection-diffusion equations with radiative source terms was developed, to model the transport of heat in the gas, in the rods, and between gas and rods. The gas flows out of the furnace while the rods flow into the furnace, in a counterflow system.

We considered a simple Newton heating/cooling term which models the heating of the rods (and the resultant gas cooling) due to convection from the hot gas to the cooler rods. We also investigated the inclusion of a radiative heating term for the rods. This term is mainly driven by radiative exposure to the furnace interior and possibly some radiative heating from the vestibule lining.

5.2 Solutions and Simulations

Various team members worked on convection models with Newton heating/cooling and radiation. In brief the approaches were

- find an analytic solution under simplified assumptions,
- quantify the geometric effects on heat transfer by radiation, and
- numerical solutions of the model.

These approaches all led to results which were largely in qualitative agreement. These are currently being compiled to give a coherent picture of the modelling held over the meeting. The preliminary principles for the vestibule design were as follows.

1. Heat loss through exhaust gas that cannot help preheat rods is inefficient, therefore the current practice of exhaust gas release through a separate flue in the furnace roof should be avoided as far as possible.
2. Heating efficiency can be increased by maximising the transfer of heat from the gas to the rods in the vestibule. Heating of rods in the vestibule by radiation from the furnace is then best avoided, to improve the possibility of extracting heat from the exhausting gases. This

indicates that it is desirable to shield the rods from oven radiation while they are in the vestibule.

3. The gas velocity and vestibule height should be selected to encourage turbulent mixing of the vestibule gas — this will facilitate efficient heating of the rods.

4. In theory the longer the vestibule the higher the final temperature of the rod entering the furnace, and hence the more efficient the heating process is. However, our simulations indicate that there is a practical limit. Improvements on heat extracted from exhaust gases begin to diminish once the vestibule gets beyond 5m in length.

Acknowledgements The moderators thank the support of the industry representative, Simon Crane, for time working with the group and organising a visit to the factory. Apart from the moderators already named at the start of the report the group included: Joshua Connor, Sam Irvine, David Jenkins, Ian Taggart, David Pontin, Raju Chowdhury, Catherine Sweatman, Raquel Gonzalez, Lindon Roberts, Kristian Kiradjiev, Edward Bissaker, and Pierluigi Cesana.

Author addresses

1. **Natalie Thamwattana**, University of Newcastle, Australia.
`mailto:Natalie.Thamwattana@newcastle.edu.au`
orcid:0000-0001-9885-3287

2. **Michael H. Meylan**, University of Newcastle, Australia.
`mailto:mike.meylan@newcastle.edu.au`
orcid:0000-0002-3164-1367

3. **A.J. Roberts**, University of Adelaide, Australia.
`mailto:ProfAJRoberts@protonmail.com`
orcid:0000-0001-8930-1552