



Monday 28 November

OSHA Healthcare Modelling Cluster

09:30 – 9:40: Opening Address

09:40 – 10:40: Plenary

- ▷ Melanie Reuter-Oppermann – Integrated Planning in Healthcare

10:40 – 11:00: Morning Tea

11:00 – 12:50: Healthcare Modelling

- ▷ Thomas Adams – Scheduling surgeries using optimisation algorithms
- ▷ Mazyar Zarepour – A districting approach to investigate the effect of home carers on the number of caregivers in home healthcare
- ▷ Michael O'Sullivan – How modelling and digital transformation can improve healthcare
- ▷ Panel Discussion – Thomas Adams, Pieta Brown, Isaac Cleland, Melanie Reuter-Oppermann, Cameron Walker, Ilze Ziedins.

12:50 – 14:40: Lunch + ORSNZ Annual General Meeting

14:40 – 15:40: Network Models & Data Science

- ▷ Golnar Behzadi – Product categorisation and implementation challenges
- ▷ Param Iyer – Is it time to retire the Solomon dataset for benchmarking VRP solution techniques?
- ▷ Karl Zhu – Finding the missing with integer programming
- ▷ Andrea Raith – Simulation-optimisation for optimal location of wireless charging infrastructure for an electric taxi fleet

15:40 – 16:00: Afternoon Tea

16:00 – 17:00: Decomposition Methods

- ▷ Ali Sohrabi – Benders decomposition algorithms to solve bi-objective linear programmes
- ▷ Michael Zhang – Speeding up column generation by heuristically limiting the solution space for airline crew rostering
- ▷ Andrew Mason – Neighbourhood pricing: A matheuristic column generation approach for staff rostering

Tuesday 29 November

09:20 – 10:40: Hydrogen Integration into Energy Systems

- ▷ Andy Philpott – On hydrogen and electricity security of supply
- ▷ Adam Clifford – Dependent Inflow models in JADE
- ▷ Uwe Langenmayr – Calculation of Synthetic Energy Carrier Production Costs with high Temporal and Geographical Resolution
- ▷ Anthony Downard – Multi-horizon modelling of hydrogen integration

10:40 – 11:00: Morning Tea

11:00 – 12:40: Production Processes & Environmental Modelling

- ▷ Ehsan Izadpanahi – Energy transition planning in manufacturing
- ▷ Geoff Leyland – Optimisation in Dairy 101
- ▷ Oleg Barbin – Using operational models to make strategic decisions
- ▷ Kerry Mayes – Multi-discipline Multi-criteria Multi-models
- ▷ Preston Ferreira – Automating the quality assurance process for continuous environmental data

12:40 – 13:40: Lunch

13:40 – 15:00: Modelling & Optimization under Uncertainty

- ▷ Dominic Keehan – Model predictive control and distributionally robust stochastic dynamic programming
- ▷ Rishi Adiga – A general framework for modelling decision-dependent information revelation in stochastic programming
- ▷ Christina Lin – Modelling housing feature impacts on sale price in newly developed suburbs relative to a standard house
- ▷ Sarah Marshall – Which flights to cancel? Disruption mitigation for a tail allocation problem

15:00 – 15:20: Afternoon Tea

15:20 – 16:40: Operations Research Tools & Supply Chain Modelling

- ▷ Maryam Mirzaei – Mapping the future of information sharing and processing in supply chains — A literature review
- ▷ Oscar Dowson – JuMP 1.0: Recent improvements to a modeling language for mathematical optimization
- ▷ Lea Kapelevich – Optimisation over exotic cones with Hypatia
- ▷ Kevin Jia – Embedding operations research and analytics into Engineering Science design

16:40 – 17:00: Prize-giving + Close

Monday 09:40 – 10:40

Plenary

Integrated Planning in Healthcare

Melanie Reuter-Oppermann. Information Systems, Technical University of Darmstadt.

Healthcare systems worldwide are complex service networks with many dependencies and interrelations. Within these networks, hospitals play a central role and serve as one of the key providers of healthcare services. Efficient and high-quality service requires the necessary resources such as personnel and beds to be dimensioned correctly and to be available in the right place at the right time. Typically, resource-related planning problems are treated separately for each resource without integrated planning across several resources, leading to sub-optimal solutions. In practice, these problems are often dealt with by clinical staff who are not trained for this kind of complex decisions and it also reduces their available time for patient treatment and care. Optimisation approaches, ideally paired with simulation models, can support decision making in healthcare practice, reduce the workload for clinical staff and decision makers and improve different objectives including efficiency, staff or patient satisfaction, quality of care or costs.

The DFG-funded early-career researchers' network InPlan-CARE aims to develop a holistic view on integrated planning problems in healthcare, particularly identifying state-of-the-art research and analysing the complexity of existing problems in hospitals. Junior researchers from Germany, Austria, Belgium, the Netherlands, New Zealand, Norway, and the United Kingdom develop advanced quantitative methods to support decision making for integrated planning in hospitals. The network aims to generate awareness for the necessity of integrated planning among hospital practitioners and within the wider research community – especially strengthening interdisciplinary research.

In a hospital, all departments are connected and every logistical decision can impact others. Due to the complexity, it is not possible to build mathematical optimisation models that address all aspects simultaneously. In addition, some decisions must be made onsite and in real-time. Simulation models can help with online decision making as well as analysing the potential consequences when transferring solutions from mathematical models into practice. With the support of the Julius von Haast Fellowship by the Royal Society of New Zealand, we aim to design an integrated simulation framework to model a complete hospital allowing to investigate all interdependencies and improvement potentials and to analyse the benefit of integrated planning.

In this talk, we want to give a motivation for integrated planning in healthcare and the hospital-wide simulation framework as well as a brief overview of the existing literature. In addition, first insights from a simulation study integrating emergency departments and (overflow) wards will be presented.

Monday 11:00 – 12:50

Healthcare Modelling

Scheduling surgeries using optimisation algorithms

Thomas Adams, Michael O'Sullivan & Cameron Walker. Engineering Science, University of Auckland, New Zealand.

This presentation will both describe the development of software for scheduling surgeries that uses an optimisation algorithm, and give a demonstration of the software. The required inputs, underlying assumptions, and current functionality of the scheduling software will be described.

The software allows for existing schedules to be visualised, lists for specific sessions to be compared, and new schedules to be created, either manually or using a scheduling algorithm. The scheduling algorithm takes into account: the patients with operations that are waiting to be performed and how long they have been waiting for; and the time within ORs allocated to specialties or surgeons in which the operations can be performed. The algorithm also uses parameters that control the trade-offs between utilisation of ORs, running overtime, and preventing patients from waiting a long time for their operations.

A districting approach to investigate the effect of home carers on the number of caregivers in home healthcare

Michael O'Sullivan, Cameron Walker & Mazyar Zarepour. Engineering Science, University of Auckland, New Zealand.

In this paper, we develop a model which focuses on districting and allocation of caregivers to districts, i.e., decisions at strategic and tactical levels. The objective is to form compact districts while maintaining workload balance. Also, the model assembles teams of caregivers with different skills to be assigned to each district. The behaviour of the proposed model is then analysed by running it using different sets of randomly generated data. We aim to examine how the number of home carers can affect the number of caregivers assigned to the districts.

How modelling and digital transformation can improve healthcare

Thomas Adams[†], Michael O'Sullivan[†], Cameron Walker[†] & Ilze Ziedins[‡]. [†]Engineering Science, University of Auckland; [‡]Statistics, University of Auckland.

The combination of modelling (including artificial intelligence) and digital transformation provides opportunities for improving the delivery of healthcare. This talk will envision a healthcare system after digital transformation and explain how existing research provides the precursors for the next generation of digital health.

Panel Discussion: Healthcare Modelling

Thomas Adams[†], Michael O'Sullivan[†], Cameron Walker[†], Ilze Ziedins[‡], Pieta Brown[‡] & Isaac Cleland[‡]. [†]Engineering Science, University of Auckland; [‡]Statistics, University of Auckland; [‡]Orion Health; [‡]RosterLab

Monday 14:40 – 15:40

Network Models & Data Science

Product categorisation and implementation challenges

Golnar Behzadi, Hamid Moghadam & Caroline Izzard. Plexure.

Product categorisation, based on product similarity, plays a critical role in product-centric digital marketing or e-commerce, and has several applications such as: price estimation of a new product, demand forecasting, similarity-based recommendation, and getting comparative insights from similar products.

The scope of this study is to examine the product similarities, while getting a lower-dimensional vector representation of each product using natural language processing. Having such a numerical representation allows us to test some measure of similarity such as cosine similarity between different products. Building product similarity solutions supports product matching requirements in digital marketing companies, activities such as product categorisation into relevant “groups” or “similar products” based on their attributes. Extensive studies already exist on this topic that capture the modelling-side of this problem; however, we are focused on implementation challenges and solutions for real-world live big data applications. Moreover, we tackle this problem in a generalisable way to adopt in a live production environment to support implementation over multiple businesses.

Is it time to retire the Solomon dataset for benchmarking VRP solution techniques?

Param Iyer & Lahiru Gunasekara. Information Systems and Operations Management, University of Auckland, New Zealand.

Vehicle routing problems with time windows (VRPTW) have appeared in peer-reviewed literature since the early 1980s. However, Solomon (1987) remains the most frequently cited article on this topic, mainly due to the wide range of problem sets, albeit synthetic, provided by this paper. Research publications regularly use these problem sets to validate new and improved heuristics and optimisation techniques in vehicle routing and its variants. However, these datasets, assuming demand distribution over a 100x100 grid with Euclidean distances, are far detached from real-world scenarios.

Since typical usage of VRP use cases focuses on urban deliveries, we propose an approach that uses standard urban layouts and population density-driven demand patterns. We use image clustering of urban maps to identify the most common designs and suggest a demand distribution accounting for the tapering of the inhabitants from city centres to suburban areas. We further recommend the circuitry factor applicable to each dataset to address the average overhead in travel paths compared to Euclidean distances. Our sample datasets allow academics and practitioners to tailor their algorithms to specific city layouts and benchmark their solution performance to meet the growing need for urban deliveries.

Finding the missing with integer programming

Karl Zhu, Andrew Mason & Anthony Downward. Engineering Science, University of Auckland.

Entity resolution (ER) — the task of determining whether multiple document records refer to the same person — is a ubiquitous problem in data analytics. Traditional ER approaches treat all matching decisions independently and only compare attribute similarities between reference pairs. However, this independence assumption omits valuable relational information in situations where the references describe a rich network of relationships between people. Collective ER — where entities are resolved jointly — incorporates this previously ignored information and can predict matches with greater accuracy. This paper describes a real-world problem faced by Parininihi ki Waitōtara (PKW) — a Māori organisation with thousands of missing shareholders — and demonstrates how collective ER can be used to identify their connections. We use Markov logic networks (MLN) to convert our domain knowledge and evidence data into a Markov network. Prediction is made by performing the most probable explanation (MPE) inference on the network, which is equivalent to a Weighted Partial MaxSAT problem in an MLN. MaxSAT problems can be formulated as an integer program (IP) and solved by a mixed IP solver. We perform experiments to empirically demonstrate how the collective ER approach is able to capture the extra relational information in an example PKW problem.

Simulation-optimisation for optimal location of wireless charging infrastructure for an electric taxi fleet

Andrea Raith. Engineering Science, University of Auckland, New Zealand.

Electric vehicles (EVs) are playing a major role in the decarbonisation of transport systems, especially in countries where electricity can be generated sustainably. Commercial vehicle fleets such as taxis are often early adopters of efficient and cost-effective vehicle technologies such as EVs. However, taxi drivers may have concerns due to limited battery range, limited availability of charging locations, inability to operate while charging, as well as having to reject passengers due to insufficient battery state-of-charge (SOC). Some of these concerns could be overcome by convenient wireless charging infrastructure that is placed so that taxis can charge while they wait for their next passenger. We describe the development of a simulation of a taxi service that operates with EV taxis. This simulation enables us to capture how an electric taxi service could operate in the future, and will help better design and locate charging infrastructure. Our simulation captures battery SOC during taxi operations and charging, and is able to consider plug-in and wireless charging. A local search heuristic is developed that identifies good locations of charging stations, and the number and type of chargers needed. We will present a case study to demonstrate the simulation-optimisation approach for charger locations of an EV taxi service, and will discuss our findings with particular focus on performance of the system with plug-in or wireless charging options.

Monday 16:00 – 17:00

Decomposition Methods

Benders decomposition algorithms to solve bi-objective linear programmes

Ali Sohrabi[†], Andrea Raith[†] & Richard Martin Lusby[‡].
[†]Engineering Science, University of Auckland; [‡]Department of Technology, Management and Economics, Technical University of Denmark.

We present three variants of Benders decomposition algorithms to solve Bi-objective linear programmes (BLPs). In these algorithms, Benders decomposition is incorporated within the bi-objective simplex algorithm by decomposing the problem into a bi-objective master problem (BM) and a bi-objective sub-problem (BS). Like the bi-objective simplex algorithm, bi-objective Benders simplex algorithm (BBSA) aims to find a set of extreme efficient solutions by proceeding iteratively from the optimiser of one objective to that of the other. In this algorithm, iteratively, an efficient solution of BM is chosen to be solved with BS, this is called exploring an efficient solution in BS. If BS is infeasible, a feasibility cut is added to BM; otherwise, an optimality cut is generated for each efficient solution of BS and added to BM. The BBSA stops when no new non-dominated point in BM needs to be explored. Since solving BS is a time-consuming step in BBSA, we propose a variant of BBSA where at most, two weighted optimality cuts are generated from BS. Another variant of the BBSA proceeds in a bidirectional fashion, where its idea is to solve the problem by simultaneously starting from the minimisers of the first and second objectives. The algorithm proceeds similarly to the second algorithm in two directions until the same non-dominated point is found from each direction. We compare and present the performance of the proposed algorithms in different bi-objective linear optimisation problems.

The **Young Practitioner Prize** is jointly sponsored by: [Optimity](#) and the [Supply Chain Company](#).



Speeding up column generation by heuristically limiting the solution space for the airline crew rostering problem

Michael Zhang[†], Andrea Raith[†], Andrew Mason[†] & Oliver Weide[‡]. [†]Engineering Science, University of Auckland; [‡]CAE.

The airline crew rostering problem involves scheduling tasks for each employee. Tasks may include flight pairings (sequence of flights that start and end at the same crew base), standby and reserve, with rest periods and days off in between. Difficulties in creating a suitable roster arise from deciding between all the possible task-employee combinations, which must also respect the airline industry's many rules, regulations, and self-imposed restrictions. We can model the crew rostering problem as a set partitioning problem. Due to the large formulations that arise, we solve the crew scheduling problem with column generation, commonly solved as resource constrained shortest path problems (RCSP). Even with column generation, it is a difficult problem to solve, as finding the optimal solution is intractable due to the exponential nature of the problem. So, we must apply acceleration techniques and heuristics to obtain a good solution in a reasonable time frame.

This presentation will introduce airline crew rostering, outline column generation and the difficulties in finding good solutions. Furthermore, it will present a strategy to artificially restrict the number of tasks seen by each RCSP to reduce the size of the problem and, with it, the computational burden on the overall solution process. We also highlight these effects when applied to real-world crew rostering problem instances.

Neighbourhood pricing: A matheuristic column generation approach for staff rostering

Isaac Cleland, Andrew Mason & Michael O'Sullivan. Engineering Science, University of Auckland, New Zealand.

Generalised set partitioning integer-programming formulations can model a wide range of problems. The columns in these formulations typically represent a subset of entities such as shifts, flights or locations. The optimisation seeks the best subset of columns that together include every entity or satisfy some generalised version of this requirement. These models can have excessively large numbers of columns, and so are often made solvable by using heuristics to select a subset of promising columns.

Columnwise neighbourhood search is a matheuristic that solves a sequence of smaller set partitioning problems in which the solution of one problem is used to create new columns that are added to the model for the next solve. These new columns are created by applying a neighbourhood rule to each column selected in the previous solution. A simple implementation of this approach showed its effectiveness by allowing us to achieve 3rd place in the 2019 VeRoLog Solver Challenge.

A more sophisticated solution approach is to create columns dynamically during the linear programme solution process using column generation, resulting in a *branch-and-price* algorithm. The ideas of neighbourhood search can be applied within branch and price by modifying the generator to solve a *neighbourhood pricing* problem where we seek new columns that lie within some neighbourhood of an incumbent column. We demonstrate this approach by testing it within software developed for rostering problems.

Tuesday 09:20 – 10:40

H₂ Integration into Energy Systems

On hydrogen and electricity security of supply

Andy Philpott[†], **Anthony Downward[‡]**, **Connor Roulston[‡]** & **Harry Thurman[‡]**. [†]*Electric Power Optimization Centre, University of Auckland*; [‡]*Engineering Science, University of Auckland*.

The New Zealand Battery Project is a Government initiative to investigate mechanisms for providing security of electricity supply for New Zealand when electricity generation is close to 100% renewable and demand for electricity has grown to enable decarbonization of the economy. Lake Onslow is one such mechanism. An alternative proposal (Southern Green Hydrogen) has been mooted by Contact Energy and Meridian Energy, who have modelled a system with storage and export of hydrogen from the lower South Island. When electricity is plentiful, e.g., from excess wind generation, hydrogen can be produced and stored, and when electricity is scarce, e.g. in a dry winter, the hydrogen plant can be shut to reduce demand.

We use the JADE hydro-thermal optimization package applied to an 11-node system to investigate the expected social costs and benefits of using electricity generated from different levels of wind generation to produce hydrogen. The model can be used to investigate the locations of new wind farms and hydrogen plants, and the effects of transmission congestion and uncertainty in reservoir inflows.

Dependent Inflow models in JADE

Adam Clifford[†], **Andy Philpott[‡]** & **Anthony Downward[†]**. [†]*Engineering Science, University of Auckland*; [‡]*Electric Power Optimization Centre, University of Auckland*.

The JADE software package is written in Julia and utilises SDDP.jl to model the New Zealand electricity system. JADE aims to find the optimal hydro-thermal generation policy for each week over a given time horizon. However, JADE currently assumes that inflows into the system are stagewise independent. This does not reflect the persistence of weather patterns, making the implementation of dependent inflow models in JADE of interest.

We model the dependence of inflows using an autoregressive model. Only the aggregate of inflows into six of New Zealand's largest reservoirs are modelled. We also trial discretisation methods to approximate the autoregressive errors and investigate their impact on the resulting policies. Following the modification of JADE, the resulting models are trained and then simulated on out-of-sample historical inflows to evaluate performance.

Based on chosen sample years, the dependent inflow model implemented in JADE finds cheaper policies in dry and wet years. However, the existing JADE implementation finds a cheaper policy in a dry year. In all years tested, the total water stored week-to-week using dependent inflows is found to be less than under independence; water is used more aggressively. Each of the error representations trialled performs similarly. These results demonstrate the potential of incorporating dependent inflow models in JADE. However, the results are yet to be vigorously verified.

Calculation of Synthetic Energy Carrier Production Costs with high Temporal and Geographical Resolution

Uwe Langenmayr. *Institute for Industrial Production, Karlsruhe Institute of Technology*.

While the transition of the electricity sector proceeds, other sectors such as industry, transportation and agriculture may fall behind. One reason is the inability to electrify all processes. Power-to-X applications allow the transformation of these sectors by replacing conventional energy carriers with renewable synthetic energy carriers.

In this approach, the production costs of different synthetic energy carriers with a high temporal and spatial resolution are calculated. Hourly weather data of Australia and New Zealand with a spatial resolution of 0.25° × 0.25° are processed into capacity profiles. These capacity profiles, covering 11 years, are clustered into profiles including the representative weeks for each cell in the covered area using k-means clustering. The production processes of green hydrogen, ammonia, methanol and green crude are modeled with a generic linear program.

The results show that especially in Australia low production costs can be achieved. In combination with large land availability, this enables large-scale synthetic energy carrier production and possible export. Hydrogen derivatives are more expensive in production but transportation might play a significant role when deciding which synthetic energy carrier should be produced. Due to the spatial proximity to Australia, New Zealand might use this circumstance to import synthetic energy carriers in the early years and start domestic production when electrolysis achieves maturity.

Multi-horizon modelling of hydrogen integration

Anthony Downward[†], **Andy Philpott[‡]** & **Evelyn Hunsberger[‡]**. [†]*Engineering Science, University of Auckland*; [‡]*Electric Power Optimization Centre, University of Auckland*; [‡]*Vector*.

Energy systems world wide are undergoing transitions to renewable sources of energy. In New Zealand, the government's goal has been to reach a 100% renewable electricity system by 2030 (however, this target has been relaxed recently). This target is challenging since increased wind capacity will lead high levels of volatility in the supply of electricity.

One way of balancing supply and demand is by building flexible hydrogen production plants, creating a renewable fuel for other sectors of the economy. These plants may be able to respond to short- and medium-term imbalances in the electricity sector.

In this talk, we will present our EMERALD model that uses the JuDGE package for Julia to create a stochastic multi-horizon model of the short- and medium-term operations of the electricity sector. This is optimised over a scenario tree, along with investments in new electricity generation and hydrogen production capacity.

We will explore the effects of different objective functions on the optimal investment policies, including an expected cost objective, as well as ones incorporating risk aversion.

Tuesday 11:00 – 12:40

Production Processes & Environmental Modelling

Energy transition planning in manufacturing: An integrated framework of machine learning and mathematical programming

Ehsan Izadpanahi & Anthony Downward. Engineering Science, University of Auckland.

Replacement of fossil fuels with clean energy sources is an important challenge for various industries that may require structural changes in manufacturers' energy systems. In some industries — e.g. dairy production — heat and electricity play a vital role in production processing and greatly affect the operations decisions. Manufacturing companies therefore need an efficient and accurate plan when executing the transition. To achieve this, mathematical modelling can be used to optimize both energy-related investment decisions as well as operations decisions, utilising machine learning algorithms as an important tool to analyse data and predict critical parameters.

In this work, two optimization models are developed: one addresses the changes in the capacity of the manufacturing firm, taking into account the environmental and financial issues; and the other considers the details of the energy system. Machine learning algorithms are developed to forecast energy consumption in terms of both electricity and heat in manufacturing and generate a scenario set. Given these scenarios, the feasibility of the energy transition plan is an important priority for manufacturing companies, a worst-case minimization approach is developed to find a robust solution. The results show that the proposed approach yields high-quality solutions that are feasible across the set of scenarios.

Optimisation in Dairy 101

Geoff Leyland. Fonterra Co-operative Group.

The dairy industry has a set of specific features & quirks that not only result in unique business practices, but also define how we model dairy optimisation problems. In this talk, I'll introduce Fonterra, the unique characteristics of modelling in the dairy industry and the challenges presented by these; some of which are still posing a problem today.

Using operational models to make strategic decisions

Oleg Barbin. Fonterra Co-operative Group.

Optimisation models can be a powerful tool in enabling strategic decision making through scenario modelling. In particular, using models that have been verified, tested & refined through operational use yields insightful and accurate results. This talk covers the lessons our team at Fonterra learned from adapting optimisation tools originally designed for "real" use to model various future scenarios and extract additional value from the models.

Multi-discipline Multi-criteria Multi-models

*Kerry Mayes[†], Bethana Jackson[#] & Electra Kalaugher[†].
[†]WSP; [#]Nature Braid.*

I will present lessons from recent work in modelling land catchments for policy development. At WSP, we have been working on projects modelling the impacts of agricultural policy at a catchment scale. We are using multiple disciplines, soil and water sciences, farm management, and economic modelling. We have multiple criteria, the impacts on the regional economy (and farmers in particular), greenhouse gases, biodiversity, eutrophication, and sediment / erosion effects. We are using multiple models, a physical model of water, contaminants, and sediment movement through the landscape (Nature Braid), models of typical farms (Farmax or custom built) and "Economic Modelling" to estimate land use change and glue the other models together. Typical challenges are for the different disciplines to understand the other disciplines well enough to integrate the solutions, inputs for one can be interim results from another – it is never as simple as you expected at the outset. Following a "Minimum Viable Product" approach we have a working process to provide sensible consistent scenario analysis. For one project all that was required was to show the benefit of this type of modelling approach, another will be used to support policy decision making. So far, no decision analytics have been required; the engineers' approach of "model it until the answer is obvious" has been working. Future work is expected to become more complicated and amenable to additional analytics.

Automating the quality assurance process for continuous environmental data

Preston Ferreira[†], Sagar Soni[#] & Kurt Joy[#]. [†]Northland Regional Council; [#]Orbica.

It is important to monitor and manage New Zealand's limited water resources given the current climate change emphasis. The Resource Management Act of 1991 stipulates that this responsibility is placed on councils to collect, process and store environmental data. This data should be timeously available to create fit-for-purpose policies; make data driven decisions; and adequately manage water resources.

Manual data processing leads to a high processing time for reliable data to be available. This is causing a growing backlog of unprocessed data, with up to 15 months backlog. This will worsen as more and more data will be collected in the future.

As a solution, Northland Regional Council (NRC) has started a project to automate the handling and cleaning of continuous environmental data to an agreed standard. The National Environmental Monitoring Standards (NEMS) will be used to ensure consistency across councils for most datasets.

Although this project was initiated by NRC, it has attracted interest from the sector as this project seeks to address common issues. Namely, inconsistencies across organisations and individuals; the process is highly resource intensive and inefficient. The project initially sets out to automate processing stage data (water level), as this is one of the larger data sets with about 5.5 million data points collected per annum across 80 sites. This presentation will share this journey including the lessons learnt and next steps.

Tuesday 13:40 – 15:00

Modelling & Optimization under Uncertainty

Model predictive control and distributionally robust stochastic dynamic programming

*Dominic Keehan[†], Andy Philpott[‡] & Eddie Anderson[‡].
[†]Engineering Science, University of Auckland; [‡]Electric Power Optimization Centre, University of Auckland; [‡]University of Sydney and Imperial College London.*

Stochastic dynamic programming problems are difficult to solve to optimality. Model predictive control (where a simpler deterministic problem is solved as an alternative) has enjoyed a wide range of application for this reason. We study an instance of model predictive control where each random value is held at its expectation. Though surprising, this approach yields policies that perform competitively in some applications. For example, Fonterra Co-operative Group Limited utilise this control policy to govern dairy product sales driven by periodic price fluctuations. Even more surprisingly, when solving stochastic dynamic programs using data-driven approximations of the true random dynamics, model predictive control can perform better out-of-sample. We study this phenomenon in an idealised version of Fonterra's application. We observe that model predictive control can be interpreted as distributionally robust stochastic dynamic programming. After demonstrating out-of-sample value in its own right, we further hypothesise that this distributional robustness explains the strong performance of model predictive control and support this with numerical studies.

A general framework for modelling decision-dependent information revelation in stochastic programming

*Rishi Adiga[†] & Andy Philpott[‡].
[†]Engineering Science, University of Auckland; [‡]Electric Power Optimization Centre, University of Auckland.*

The stochastic programming framework for optimising decision-making assumes a predetermined information structure, where certain information is inevitably revealed at fixed time stages regardless of the decisions made. This assumption is not valid for problems with endogenous uncertainty, in which the sequence of information revelation is not fixed beforehand and is instead governed by model decisions. In this paper we present a comprehensive framework for defining information structures for such programs, based on modelling decision-dependent information discovery in an exploration-exploitation paradigm. This enables modellers to keep track of information flows and ensure that the desired causality between model decisions and information revelation is modelled appropriately. This framework is applied on a SMIP optimising well placement for a geothermal reservoir, and results demonstrating the effect of the information structure are discussed.

Modelling housing feature impacts on sale price in newly developed suburbs relative to a standard house

*Christina Yin-Chieh Lin[†], Andrew Mason[†], Charles Ma[‡] & Andreas W. Kempa-Liehr[†].
[†]Engineering Science, University of Auckland; [‡]MADE Group Limited.*

There is a recent trend of entire new suburbs that support a local community being designed and built to solve the shortage of affordable housing all around the world. The aim of this study is to anticipate the value of housing features in new suburbs that are still in the planning stage. For this purpose, we are separating price movements over time from the impact of individual housing features in recently developed suburbs. To generate insights on housing features that can be directly interpreted by developers, we propose modelling house prices relative to a standard house representative of local preferences. The proposed model is successfully evaluated on newly developed suburbs in Auckland, New Zealand. The case study on the newly developed suburbs of Fairview Heights, Oteha, Stonefields, Long Bay, Flat Bush, and Silverdale demonstrates that the proposed modelling approach effectively captured the complex relationship between housing features and sale price relative to a standard house (R-squared 91.5%). The proposed model generalizes to a reasonable extent to house prices in the new suburb of Hobsonville (R-squared 75.0%) without using any historical sale records in this suburb. This indicates that the insights on housing features relative to the standard house are applicable to other new suburbs still in the planning stage and, therefore, have the potential to support future suburb developments.

Which flights to cancel? Disruption mitigation for a tail allocation problem

Sarah Marshall. Department of Mathematical Sciences, Auckland University of Technology, New Zealand.

Tail allocation describes the process of assigning scheduled flights to a particular aircraft (a tail). This optimization problem can be formulated as a mixed-integer linear programming problem. However, when unexpected disruptions occur (e.g. due to crew illness, or bad weather), the number of available aircraft is often reduced, meaning that not all flights can be flown. The question then becomes, which flights should be cancelled in order to minimize the disruption. This is a challenging problem as there are multiple competing objectives. This research explores a variety of disruption mitigation strategies and investigates their impact on several performance measures.

Tuesday 15:20 – 16:40

Operations Research Tools & Supply Chain Modelling

Mapping the future of information sharing and processing in supply chains — A literature review

Maryam Mirzaei & David Riverton. Business Information Systems, Auckland University of Technology, New Zealand.

With continuing, digitalisation of operations and the adoption of new technologies, such as the IoT, big data, and blockchain, the nature of Information sharing and processing will be transformed in the near future, which in turn will transform the supply chains.

Given that supply chains are complex and interconnected, these changes require substantial investments and will impact existing processes, organisational structures, and roles. This presentation provides an overview of a comprehensive literature review that attempts to understand characteristics that differentiate Information sharing and processing in the current supply chain from the future supply chain.

A total of 31 peer-reviewed journal articles were identified and synthesised. Papers were selected and classified based on emerging and future trends and innovations as well as step changes and disruptions impacting information sharing and processing in supply chain management.

The presentation will provide an outline of scholarly views on the dominant features of future supply chains as well as key technologies that will transform supply chains.

JuMP 1.0: Recent improvements to a modeling language for mathematical optimization

Oscar Dowson. JuMP.

JuMP is an algebraic modeling language embedded in the Julia programming language. JuMP allows users to model optimization problems of a variety of kinds, including linear, integer, conic, semidefinite, and nonlinear programs, and handles the low-level details of communicating with solvers.

After nearly 10 years in development, JuMP 1.0 was released in March, 2022. In this talk we demonstrate the unique features of JuMP that make it an ideal tool for the OR practitioner, whether they are teaching new students with minimal programming experience or solving large scale industrial nonlinear programs with millions of variables and constraints.

Optimisation over exotic cones with Hypatia

*Lea Kapelevich[†], Chris Coey[#] & Juan Pablo Vielma[‡].
[†]Mathworks; [#]Operations Research Center, MIT, Cambridge, MA.; [‡]Google Research and MIT Sloan School of Management, Cambridge, MA.*

Any convex optimisation problem can be posed as a minimisation of a linear function over the intersection of a linear subspace and a cone. We will introduce some “exotic cones” and the interior-point solver Hypatia, which allows users to define their own exotic cone types. We will cover algorithmic enhancements that make Hypatia a performant solver and give examples of efficient oracles for useful (or interesting) exotic cones.

Embedding operations research and analytics into Engineering Science design

Kevin Jia. Engineering Science, University of Auckland.

The second-year Engineering Science Design course, ENGSCI 263, underwent a significant re-design starting in 2019 as part of the Engineering Science curriculum review. Half of the course is now dedicated to an operations research and analytics project, where students develop a realistic-sized operations research model on a vehicle routing problem and simulate the resulting solution to get a measure of solution quality. The project also develops students’ professional skills as demanded by industry, such as writing, spreadsheeting. Consideration is also given to systems design concepts and data ethics.

The lessons learned after four years of teaching this course (including two years of lockdown teaching) will be discussed, and potential ideas for further course and curriculum development explored.