

ORSNZ Annual Conference 2026

Operations Research Society of New Zealand (ORSNZ)

November 19-20, 2026

The University of Auckland, Auckland, New Zealand

Opening

Whakataka te hau ki te uru	Cease the winds from the west
Whakataka te hau ki te tonga	Cease the winds from the south
Kia mākinakina ki uta	Let the breeze blow over the land
Kia mātaratara ki tai	Let the breeze blow over the ocean
E hī ake ana te atakura	Let the red-tipped dawn come with a sharpened air
He tio, he huka, he hau hū	A touch of frost, a promise of a glorious day.

Tīhei mauri ora!

Welcome

Tēnā tatou e hoa ma	Greetings to us all
Kua haere mai nei Ki te ako I ngā whakaaro	Who have come here To learn The ideas/concepts/thoughts/views
Ka mihi ki te iwi kāinga, ki a Ngāti Whātua Ōrākei	We acknowledge and give thanks to Ngāti Whātua Ōrākei as the people on whose lands we are meeting
Kānui te mihi atu Mō tō koutou aronui Ki te aro mai I ngā kaupapa	Great is the thanks/appreciation For your interest To attend The programme
Otirā, kua mutu i konei ā ku mihi mō tēnei wā Nō reira, tēnā koutou, tēnā koutou, ā, tēnā tātou katoa.	Thus, I end my greeting here for the present. Therefore, greetings and honour to one and all.

Contents

Opening **i**

Welcome **ii**

Delegate Information **1**

Programme Summary **2**

Thursday 19th November 2026 **5**

Session 1: Healthcare – 9:20am-10:20am 5

Multi-Objective Scheduling and Routing for Preventive Home Healthcare Services for Geriatric Patients (***Mohsen Abdoli***, *Sarah Marshall, Lisa Aoki Hillas*) 5

Prediction Versus Waiting as a Screening Device (***Lisa Aoki Hillas***, *Rhema Vaithianathan, Joe Steele*) 6

Revealed-Practice Scheduling: A Historically Grounded Optimization Approach for Operating Room Planning (***Signý Kristín Sigurjónsdóttir***, *Helga Ingimundardóttir, Cameron Walker, Tómas P. Rúnarsson*) 6

Session 2: Energy – 10:40am-12:00pm 7

Bridging stochastic hydro-scheduling and deterministic capacity expansion models (***Madison L. Zegeer***, *Rasul Satymov, Rebecca A.M. Peer, Jannik Haas*) 7

Optimising BESS energy arbitrage in the NZEM (***Alexandre Fargier***, *Andy Philpott*) 8

Staged capacity expansion with long-term hydro storage using cyclic SDDP policy graphs (***Ali Kök***, *Andy Philpott*) 8

Integrated Wind Farm Layout and Cable Routing: Upper Bounds and Algorithms (***Asta Susan Rustad***, *David Pisinger*) 9

Session 3 – 12:40pm-2:00pm 9

Asymptotic analysis of infinite-server queues driven by a Cox-BESQ process (***Yajun Shi***, *Jose Da Fonseca, Wenjun Zhang*) 9

A Faster Algorithm For Multi-Criteria Betweenness (***Harry Newrick***, ***Chris Oliver***, *Andrea Raith*) 10

Bustrify: A simulation and analysis package for public transport electrification (***Ahmed Abdelhadi***, *Andrea Raith, Minh Kieu*) 10

Session 4 – 2:10pm-3:30pm 11

Data-driven robust optimization for resilient transportation network planning under landslide uncertainty (***Jiaying Gu***, *Yi Wang, Seósamh B. Costello, Amelia Fabia Lin, Liam Wotherspoon*) 11

Delivery performance of platform-owned and marketplace products in a platform-operated logistics network (***Hannah Liu***) . . . 11
Planning of Centralised-Charging Battery Swapping Networks with Dynamic Service Availability (***Hanqing Wang***, *Enjian Yao, Andrea Raith, Yang Yang, Tianyu Zhang*) . . . 12
Power, Prices and Profits in Energy-Only Wholesale Electricity Markets (***Andy Philpott***) . . . 13
Session 5: Gurobi Community Event – 4:00pm-5:30pm . . . 13

Friday 20th November 2026 **14**
Session 6 – 9:00am-10:20am . . . 14
Optimising Computational Performance of Mixed Integer Programming Models (***Ava Purvis***, ***Samriddhi Jaiswal***) . . . 14
Autonomous Orchestration of Third-Party Optimisation Engines (***Maz Lakadia***) 15
Learning and Transferring Scheduling Heuristics through Genetic Programming for Dynamic Multi-Mode Projects (***Yuan Tian***, *Yi Mei, Mengjie Zhang*) 15
Recent developments in solving large linear programs (***Simon Bowly***) . . . 16
Session 7 – 10:40am-12:00pm . . . 17
Who keeps the lights on? A two-stage stochastic model of field-workforce design for electricity network resilience in Auckland (***Ehsan Mahmoodi***) . . . 17
Continuous Representations of Pareto Fronts via High-Order KKT Expansions (***Jack Yarndley***, *Roberto Armellin*) . . . 17
Column Generation Algorithm for Parallel Task Scheduling with Communication Delays (***Lewis Azzopardi***, *Emmanuel Araba, Oliver Sinnen*) . . . 18
Experiments with Neighbourhood Pricing to speed up Column Generation (***Andrew Mason***, *Basile Blayac, Andrea Raith*) . . . 18
Plenary – 1:00pm - 2:00pm . . . 19
Say What, Not How: Constraint Modelling with MiniZinc (***Jip J. Dekker***) . . 19
Session 8 – 2:20pm-3:00pm . . . 19
Spare Parts Inventory Management Utilizing Live Machine Health Data: A Digital Twin Approach (***Yale T. Herer***, *Jennifer Kruman*) . . . 19
Benefit of the Doubt Models: Input Minimisation, Zero-Valued Measures and Weight Robustness (***Paul Rouse***, *Maryam Hasannasab, Dimitris Margaritis*) 20

Closing **21**

Author Index **22**

Delegate Information

Conference Website

Detailed information about the conference is available on the conference website: `https://orsnz.org.nz/conf58/`

Venue

The conference will take place at the University of Auckland at the Science Centre (which is building 302), 23 Symonds Street. We will be in room G20 on the ground floor.

Presentations

Oral presentations will be 20 minutes in length (15 + 5 for questions).

Wifi

Wifi will be available at the conference. Details will be available at registration.

Conference Dinner

The conference dinner will be held from 5:30 pm on Thursday 19th November at The Brit Pub & Eatery (122 Quay St).

Sponsors

We are grateful for the generous sponsorship from this year's sponsors:









ORSNZ Conference Programme

Thursday 19th November 2026

8:30am–9:00am	**Registration**
9:00am–9:20am	**Opening**
9:20am–10:20am	**Session 1: Healthcare** (Chair: Cameron Walker) Mohsen Abdoli, University of Auckland, *Multi-Objective Scheduling and Routing for Preventive Home Healthcare Services for Geriatric Patients* **(YPP, JAG)** Lisa Aoki Hillas, University of Auckland, *Prediction Versus Waiting as a Screening Device* **(JAG)** Signý Kristín Sigurjónsdóttir, University of Iceland, *Revealed-Practice Scheduling: A Historically Grounded Optimization Approach for Operating Room Planning* **(JAG)**
10:20am–10:40am	*Morning Tea*
10:40am–12:00pm	**Session 2: Energy** (Chair: Andy Philpott) Madison L. Zegeer, University of Canterbury, *Bridging stochastic hydro-scheduling and deterministic capacity expansion models* **(YPP)** Alexandre Fargier, University of Auckland, *Optimising BESS energy arbitrage in the NZEM* **(YPP)** Ali Kök, TU Wien, *Staged capacity expansion with long-term hydro storage using cyclic SDDP policy graphs* **(YPP, JAG)** Asta Susan Rustad, Technical University of Denmark, *Integrated Wind Farm Layout and Cable Routing: Upper Bounds and Algorithms* **(YPP, JAG)**
12:00pm–1:00pm	*Lunch*
1:00pm–2:00pm	**Session 3** (Chair: Lisa Aoki Hillas) Yajun Shi, Auckland University of Technology, *Asymptotic analysis of infinite-server queues driven by a Cox-BESQ process* **(YPP)** Harry Newrick & Chris Oliver, University of Auckland, *A Faster Algorithm For Multi-Criteria Betweenness* **(YPP)** Ahmed Abdelhadi Ahmed Abdelhadi, University of Auckland, *Bustrify: A simulation and analysis package for public transport electrification* **(YPP, JAG)**
2:00pm–2:10pm	*Short break*
2:10pm–3:30pm	**Session 4** (Chair: Sarah Marshall) Jiaying Gu, University of Canterbury, *Data-driven robust optimization for resilient transportation network planning under landslide uncertainty* **(YPP)** Hannah Liu, University of Auckland, *Delivery performance of platform-owned and marketplace products in a platform-operated logistics network* **(JAG)** Hanqing Wang, Beijing Jiaotong University, *Planning of Centralised-Charging Battery Swapping Networks with Dynamic Service Availability* Andy Philpott, University of Auckland, *Power, Prices and Profits in Energy-Only Wholesale Electricity Markets*
3:30pm–4:00pm	*Afternoon Tea*
4:00pm	**Gurobi Community Event** Michael O'Sullivan, University of Auckland, *Optimisation in practice: water resource management planning* Hamideh Anjomshoa, Gurobi, *Gurobi Intelligence: GenAI in Gurobi* Simon Bowly, Gurobi, *What's new at Gurobi: Gurobi 14.0*
5:30pm	**Conference Dinner** The Brit Pub & Eatery, 122 Quay St.

Friday 20th November 2026

9:00am–10:20am	**Session 5** (Chair: Michael O'Sullivan) Ava Purvis & Samriddhi Jaiswal, University of Auckland, *Optimising Computational Performance of Mixed Integer Programming Models* Maz Lakadia, Evolve DS, *Autonomous Orchestration of Third-Party Optimisation Engines* Yuan Tian, Victoria University of Wellington, *Learning and Transferring Scheduling Heuristics through Genetic Programming for Dynamic Multi-Mode Projects* Simon Bowly, Gurobi, *Recent developments in solving large linear programs*
10:20am–10:40am	*Morning Tea*
10:40am–12:00pm	**Session 6** (Chair: Andrea Raith) Ehsan Mahmoodi, University of Auckland, *Who keeps the lights on? A two-stage stochastic model of field-workforce design for electricity network resilience in Auckland* Jack Yarndley, University of Auckland, *Continuous Representations of Pareto Fronts via High-Order KKT Expansions* Lewis Azzopardi, University of Auckland, *Column Generation Algorithm for Parallel Task Scheduling with Communication Delays* Andrew Mason, University of Auckland, *Experiments with Neighbourhood Pricing to speed up Column Generation*
12:00pm–12:20pm	*Lunch*
12:20pm–1:00pm	**ORSNZ Annual General Meeting**
1:00pm–2:00pm	**Plenary talk** (Chair: Andrew Mason) Jip J. Dekker, Monash University, *Say What, Not How: Constraint Modelling with MiniZinc*
2:00pm–2:20pm	*Afternoon Tea*
2:20pm–3:00pm	**Session 7** (Chair: David Robb) Yale T. Herer, Technion – Israel Institute of Technology, *Spare Parts Inventory Management Utilizing Live Machine Health Data: A Digital Twin Approach* Paul Rouse, University of Auckland, *Benefit of the Doubt Models: Input Minimisation, Zero-Valued Measures and Weight Robustness*
3:00pm	**Prizegiving and Closing**

Thursday 19th November 2026

Session 1: Healthcare – 9:20am-10:20am

Chair: Cameron Walker

Multi-Objective Scheduling and Routing for Preventive Home Healthcare Services for Geriatric Patients

19 Nov
9:20am-
9:40am

Mohsen Abdoli, Sarah Marshall, Lisa Aoki Hillas
University of Auckland

Preventive home healthcare plays an important role in supporting older adults by identifying changes in their health or condition early and reducing the need for more intensive care. However, limited caregiver capacity, travel requirements, and scheduling constraints make it difficult to deliver preventive visits regularly and as prescribed. Although existing home healthcare scheduling and routing models consider various operational objectives, the longer-term effects of missed preventive care on patient health remain relatively underexplored. This study develops a multi objective, multi-period mixed-integer linear programming model that integrates weekly patient selection, daily visit assignment, caregiver routing, and patient health dynamics. The model considers two competing objectives: minimizing cumulative deterioration burden and minimizing travel cost. Changes in patient health are represented through a time-dependent formulation that considers natural deterioration, the effect of preventive visits, and additional deterioration associated with unmet care needs. Weighted-sum and -constraint approaches are used to examine the trade-offs between the two objectives. Preliminary computational experiments using a small illustrative instance demonstrate that the model can generate feasible schedules and routes while showing how service decisions affect patient health trajectories over time. The results also reveal a trade-off between reducing deterioration burden and improving travel efficiency. Overall, this research provides an initial framework for connecting short-term operational planning decisions with longer-term patient health considerations in preventive home healthcare.

Prediction Versus Waiting as a Screening Device

19 Nov
9:40am-
10:00am

Lisa Aoki Hillas, Rhema Vaithianathan, Joe Steele
University of Auckland

Rough estimates indicate that on a given night in the US, about there are about 400,000 individuals in homeless shelters and about 200,000 individuals sleeping on the streets. US federal rules prioritise scarce permanent supportive housing for people who are "chronically homeless" (homeless for at least 12 months) The rationale is screening: many people exit homelessness quickly without help, so requiring waits intends to allocate housing to those most in need. But waiting is costly, because the people at highest risk can die while they wait. Predictive risk models offer an imperfect way of classifying individuals when they are first seen in a shelter.

We model housing allocation as a scheduling problem with two types of impatient customer: high-need individuals, who leave the queue only by dying, and low-need individuals, who leave by resolving their own homelessness. We assume a planner who wants to minimise deaths before housing. Using a fluid model following Bassamboo, Randhawa and Wu (2023) and simulations calibrated to San Francisco data, we consider when waiting time and a classifier act as substitutes or complements. Preliminary results suggest that in a congested system the chronicity rule does not bind without a classifier. With an imperfect classifier, the rule can either reduce deaths, by filtering out misclassified low-need people who self-resolve, or increase them, by delaying housing for high-need people; which effect dominates depends on how mortality and self-resolution risks evolve over time spent homeless. We quantify how the rule's value changes with classifier accuracy and housing supply.

Revealed-Practice Scheduling: A Historically Grounded Optimization Approach for Operating Room Planning

19 Nov
10:00am-
10:20am

Signý Kristín Sigurjónsdóttir, Helga Ingimundardóttir, Cameron Walker, Tómas P. Rúnarsson
University of Iceland

Operating room scheduling models are rarely used in practice, partly because they do not match how hospitals schedule. We propose revealed-practice scheduling, where historical schedules are used to estimate parameters and to limit where the optimizer searches. From waiting-list snapshots at the National University Hospital of Iceland, we learn operation windows, room-day surgery patterns, the desirability of surgeon-room-weekday combinations, and per-surgeon daily workload limits. A mixed-integer program assigns surgeries to days and rooms across nine orthopaedic operating rooms, subject to these constraints and ICU and ward capacity. Its lexicographic objective first minimizes unscheduled surgeries, then moves of planned surgeries, and then maximizes desirability. The learned constraints describe typical rather than exceptional practice, so real schedules will sometimes break them. The current plan is always feasible, and exceptions are left to the scheduler.

We evaluate on 89 weekly planning dates. For each date, we rebuild the schedule as it stood, give the model the same surgeries the hospital faced, without revealing where the hospital later placed them, and compare its schedule with what the hospital did. First, we

check how often the real schedules break each learned constraint. Violations split evenly between patterns, workload limits and operation windows, and show both where the model is stricter than practice and where the hospital went beyond its usual limits. Second, we compare the two schedules. The model schedules 96% of what the hospital performed or kept, uses a similar number of moves, and places surgeries in more desirable slots every week.

Session 2: Energy – 10:40am-12:00pm

Chair: Andy Philpott

Bridging stochastic hydro-scheduling and deterministic capacity expansion models

19 Nov
10:40am-
11:00am

Madison L. Zegeer, Rasul Satymov, Rebecca A.M. Peer, Jannik Haas

Sustainable Energy Research Group, Department of Civil and Environmental Engineering, University of Canterbury

The global transition to a sustainable energy system requires not only ambitious policy targets but also modelling tools that can translate decarbonisation goals into investment and operational decisions. In hydro-dominated electricity systems, this task falls between two modelling traditions. Stochastic hydro-scheduling methods, such as Stochastic Dual Dynamic Programming (SDDP), represent reservoir operation under inflow uncertainty, whereas capacity expansion models, which inform long-term investment, typically rely on deterministic weather and inflow assumptions.

In Aotearoa New Zealand's isolated electricity system, hydropower provides ˜60% of annual supply and most of the long-duration energy storage. This storage is limited relative to demand, cycled within a year, and concentrated in a few catchments, leaving the system vulnerable to dry years. Investment decisions based on deterministic inflows may therefore underestimate the probability and cost of supply shortfalls.

This study extends our previous work that validated reduced-order stochastic hydro formulations against SDDP. We are incorporating inter-seasonal storage constraints derived from stochastic hydro-scheduling into the capacity expansion model PyPSA-NZ. The framework will evaluate transition pathways to 2050 across weather years, demand trajectories and reliability options.

Preliminary deterministic dry-year results show that wind and solar generation displaces hydro dispatch through summer and autumn, conserving reservoir storage ahead of winter. However, these runs assume perfect foresight of inflows and permit aggressive drawdown, which likely overstates the firming value of this seasonal complementarity. Ongoing work applies the stochastic constraints to test whether this result holds when reservoir operation must hedge against inflow uncertainty.

Optimising BESS energy arbitrage in the NZEM

19 Nov
11:00am-
11:20am

Alexandre Fargier, Andy Philpott
University of Auckland

In recent years, Battery Energy Storage Systems (BESS) have emerged as an important component of the New Zealand Electricity Market (NZEM). As BESS participation increases, the effectiveness of the regulatory framework governing their operation becomes increasingly important. However, the current NZEM regulatory framework requires a one-hour gate closure that presents operational challenges for BESS participation in the energy market. Under these conditions, BESS must commit to trading decisions despite significant uncertainty, which could result in conservative trading strategies and underutilisation of these assets. While proposals to shorten or eliminate gate closure are currently under consideration, this work focuses on regulatory frameworks that address this constraint without requiring changes to gate closure itself. In this study, we develop a family of dynamic programming models to optimise BESS energy arbitrage under current and alternative regulatory frameworks. Using these models, we compare the costs and benefits of different regulatory frameworks for BESS operators and assess their implications for the wider market.

Staged capacity expansion with long-term hydro storage using cyclic SDDP policy graphs

19 Nov
11:20am-
11:40am

Ali Kök, Andy Philpott
Energy System Economics Lab, Institute of Net-Zero Energy Systems, TU Wien

Hydro-dominated systems such as New Zealand's make capacity planning hard because the value of stored water depends on decisions far into the future. Recent work solves this with stochastic dual dynamic programming (SDDP) on a cyclic policy graph: a single investment decision followed by an infinite-horizon operating loop. We extend this approach to investment under uncertain demand growth. We first derive, from the discounted cost, several policy-graph designs for staged investment, in which the investment nodes enter infinite operating loops that are shared or kept separate. These designs resemble multi-horizon stochastic programs, which price each investment period with a representative operating period. The difference is that our operating periods are infinite SDDP loops, so the value of water carried from year to year is found inside the model rather than fixed by assumed start and end storage levels. Like most multi-horizon programs, however, the loops cannot carry storage from one investment decision to the next; on a hydro-thermal test system the resulting cost error grows with reservoir size. We therefore propose a design for the first decision: an investment with a lead time, a few explicit operating years that carry storage forward, a branch into demand scenarios, and one infinite operating loop per scenario. The lead time makes the first decision a commitment made before demand is known. We are applying the design to a three-region model of the New Zealand system with solar investment and demand-growth scenarios, such as new data centres, to study investment and electricity prices.

Integrated Wind Farm Layout and Cable Routing: Upper Bounds and Algorithms

19 Nov
11:40am-
12:00pm

Asta Susan Rustad, David Pisinger
Technical University of Denmark

Accelerating the green transition requires lowering the cost of sustainable energy. For offshore wind farms, costs are largely driven by two factors: the placement of the turbines and the layout of the interconnecting cable network. Because of the wake effect, where turbines "steal" wind from those downstream, turbine placement favors a dispersed layout. Cable costs, however, grow with cable length and therefore favor clustered turbines. Although each problem has been studied intensively on its own, their conflicting objectives call for a unified framework to avoid suboptimal choices.

In this presentation, we introduce upper and lower bounds on cabling costs and show how incorporating them into the turbine placement model makes it possible to find a good compromise between placement and cabling costs. We further combine these into overall upper and lower bounds for the unified layout and cabling problem. To the best of our knowledge, these are the first polynomial bounds in the literature, and we show experimentally that they are relatively tight.

Building on these results, we present a semi-integrated solution method. The placement problem is solved with a Variable Neighborhood Search, guided by a weighted sum of the upper and lower cabling-cost bounds. While the estimated cable costs thus influence the turbine placement, the final cabling layout is determined with a matheuristic.

We report computational experiments on realistic instances with more than 21,000 candidate locations and over 150 turbines.

Session 3 – 12:40pm-2:00pm

Chair: Lisa Aoki Hillas

Asymptotic analysis of infinite-server queues driven by a Cox-BESQ process

19 Nov
1:00pm-
1:20pm

Yajun Shi, Jose Da Fonseca, Wenjun Zhang
Auckland University of Technology

This paper studies asymptotic properties of infinite-server queues driven by a Cox process whose stochastic arrival intensity follows a squared Bessel (BESQ) process. Using previously derived transform representations available through the affine structure of the model, we establish laws of large numbers, central limit theorems, large deviation principles, and moderate deviation principles. These results are obtained for three important settings: cumulative occupancy over a finite time interval, system occupancy under a diverging long-run mean arrival intensity, and the superposition of independent infinite-server queues with Cox-BESQ arrival intensities. Numerical experiments illustrate the associated rate functions and demonstrate the effectiveness of the resulting rare-event approximations.

A Faster Algorithm For Multi-Criteria Betweenness

19 Nov
1:20pm-
1:40pm

Harry Newrick, **Chris Oliver**, Andrea Raith
Uinversity of Auckland

Betweenness centrality identifies important components in a network by how often they lie on shortest paths. In transport networks, routes are rarely judged on a single cost; a recent bi-criteria extension shares credit equally across all Pareto-optimal paths between each origin-destination pair, thereby capturing trade-offs between criteria such as distance and travel time. However, this removes the structure that makes Brandes' algorithm the most efficient for single-criteria betweenness computation: a vertex can be reached through multiple non-dominated cost combinations, so its paths cannot be summarised by Brandes' path counts and predecessor relationships.

We present an accumulation method that restores the spirit of Brandes' backward pass in the bi-criteria case. Rather than constructing paths, we treat the non-dominated labels produced by a one-to-all bi-criteria label-setting algorithm as the vertices of a directed acyclic graph, each linked to the labels it extended. Betweenness credit is initialised at the labels of each destination and passed backwards in reverse order of settlement. This produces a traceable topological order, so every label and predecessor link is processed once, no matter how many paths use it. Distinct paths with identical cost vectors must each be counted, so equivalent labels are merged, given a path count, and credit is split among predecessors in proportion to these counts. Accumulation runs inside the solver as each one-to-all path computation completes, so memory does not grow with the number of origins. The method reduces to classical betweenness computation for a single criterion and applies to any number of criteria.

Bustrify: A simulation and analysis package for public transport electrification

19 Nov
1:40pm-
2:00pm

Ahmed Abdelhadi, Andrea Raith, Minh Kieu
University of Auckland

Electric-bus planning requires coordinated decisions concerning service schedules, fleet operations, charging strategies, infrastructure deployment, and onboard energy-storage capacity. We present Bustrify, a Julia-based simulation-optimisation package for constructing, evaluating, and planning electrified bus networks. Bustrify processes GTFS data to reconstruct routes, generate operational timetables and bus schedules, and map transit services onto detailed road networks derived from sources such as OpenStreetMap. Its discrete-event simulation supports both preassigned bus blocks and dynamic trip dispatching, whereby scheduled trips are released over time and assigned in real-time. The framework represents repositioning movements, passenger operations, traffic conditions, charging behaviour, infrastructure utilisation, queueing, and system-level control decisions. A longitudinal vehicle dynamics energy model is developed to estimate edge-level energy expenditure using road-network elevation profiles, passenger-dependent vehicle mass, and regenerative recovery. To demonstrate the framework's versatility across transit technologies and planning workflows, an Auckland case study evaluates a 100-kW static wireless opportunity-charging

configuration under progressively reduced battery capacities to investigate the operational feasibility of this emerging technology. The case study represents 29,377 weekly trips served by a fleet of 300 buses and 66 wireless chargers distributed across 26 layover locations.

Session 4 – 2:10pm-3:30pm

Chair: Sarah Marshall

Data-driven robust optimization for resilient transportation network planning under landslide uncertainty

19 Nov
2:10pm-
2:30pm

Jiaying Gu, Yi Wang, Seósamh B. Costello, Amelia Fabia Lin, Liam Wotherspoon
University of Canterbury

Road transportation networks are essential to societal and economic functioning yet remain vulnerable to earthquake-triggered landslides. This paper develops a data-driven robust optimization model for pre-event investment planning to enhance network resilience under co-seismic landslide uncertainty. A tri-level defender–nature–traveler model is formulated to capture interactions among infrastructure reinforcement decisions, worst case hazard realizations, and traveler response. A data-driven polyhedral uncertainty set is constructed from hazard scenarios, providing a tractable and behaviorally interpretable representation of spatially correlated failures. The framework is applied to the South Island state highway network of New Zealand. Results show that reinforcement priorities are jointly shaped by hazard exposure, network topology, and demand, with landslide-prone, high-demand links emerging as the first bottlenecks to be reinforced. As budgets increase, reinforcement strategies shift beyond these bottlenecks, indicating that optimal investment patterns cannot be extrapolated from low-budget priorities. Sensitivity analysis reveals a trade-off between pre-event reinforcement and post-event restoration.

Delivery performance of platform-owned and marketplace products in a platform-operated logistics network

19 Nov
2:30pm-
2:50pm

Hannah Liu
University of Auckland

Many online retail platforms sell their own products and also host independent sellers, and some allow these sellers to use the platform's logistics network. This study examines whether orders from the two groups receive comparable delivery service when they are delivered through the platform's network. Using transaction-level data from JD.com for March 2018, we compare delivery time for first-party orders (the platform's own products) and marketplace orders (products from independent sellers). Delivery time is divided into order processing, shipping between facilities and last-mile delivery. We also examine the delivery time promised to customers and whether orders arrive by the promised deadline. Our findings suggest that marketplace orders take about one day longer to reach customers on average. Most of this difference occurs before parcels reach the delivery station, and

last-mile delivery times are similar for the two groups. Marketplace orders also receive longer promised delivery time. With these longer promise times, the share of orders that arrive by their promised date is similar for the two groups. When orders are given the same promised delivery time, however, marketplace orders arrive on time less often. This indicates that on-time rates alone may not reflect differences in delivery speed between seller types. We discuss implications for fulfilment decisions by sellers and platforms and for the measurement of delivery service.

Planning of Centralised-Charging Battery Swapping Networks with Dynamic Service Availability

19 Nov
2:50pm-
3:10pm

Hanqing Wang, Enjian Yao, Andrea Raith, Yang Yang, Tianyu Zhang
Beijing Jiaotong University

Centralised-charging battery swapping networks decouple battery charging from vehicle energy replenishment, making station service availability dependent not only on installed swapping capacity but also on the spatiotemporal matching between battery consumption and replenishment. However, existing studies often give limited attention to how battery replenishment affects users' station choices through inventory-dependent service conditions. The resulting feedback between demand assignment and replenishment decisions remains insufficiently represented in network planning. This study proposes a nested bilevel planning framework that incorporates a coupled operational equilibrium between demand assignment and battery logistics. At the upper level, the model determines battery-swapping-station locations and capacities considering construction cost, battery-delivery cost and service quality. At the lower level, a user-equilibrium demand-assignment model is coupled with a multi-depot, multi-vehicle, multi-trip vehicle routing model for battery delivery. Station service availability links the two models, enabling mutually consistent decisions on demand assignment and battery delivery. To solve the model, the Non-dominated Sorting Genetic Algorithm II (NSGA-II) is integrated with an iterative optimisation-and-assignment procedure that alternately solves the demand assignment and battery delivery problems. A case study in Tianjin is conducted to evaluate the planning and operational implications of the proposed framework. The analysis traces the evolution of battery inventories, service availability, and station queues under time-varying demand. Comparison with a decoupled planning approach further assesses the value of capturing the endogenous feedback between demand and supply in balancing planning costs and user service quality.

Power, Prices and Profits in Energy-Only Wholesale Electricity Markets

19 Nov
3:10pm-
3:30pm

Andy Philpott
University of Auckland

We review investment planning using screening curve analysis for perfectly competitive energy-only wholesale electricity markets through the lens of multistage stochastic programming. This enables new insights to be established for the optimal sizing of variable renewable energy sources and storage devices. In energy-only markets the cost of capacity is covered by economic rents earned from energy prices. A barrier to entry in any technology yields surplus rents that increase profits for that technology above long-run marginal costs, signalling that wholesale energy prices are higher than they need to be to cover capital costs. We illustrate this using a model of New Zealand's electricity system.

Session 5: Gurobi Community Event – 4:00pm-5:30pm

Friday 20th November 2026

Session 6 – 9:00am-10:20am

Chair: Michael O'Sullivan

Optimising Computational Performance of Mixed Integer Programming Models

19 Nov
9:00am-
9:20am

Ava Purvis, **Samriddhi Jaiswal**
University of Auckland

Infrastructure asset management (IAM) problems are often modelled as mixed-integer programs. However, these large-scale problems can challenge commercial solvers such as Gurobi due to memory and solve time limitations, motivating methods to improve solver efficiency without altering the underlying problem formulations.

This project investigated whether formulation structure can inform branching priorities via Gurobi's BranchPriority attribute, improving performance as measured by branch-and-bound node counts and solve times relative to Gurobi's defaults.

Bin packing problems (BPPs) served as the testbed, followed by two case studies: synthetic IAM problems, formulated as multiple-choice knapsack problems with precedence constraints, and vehicle routing problems (VRPs). BPP priority measures were based on variable constraint coefficients and constraint occurrence, alongside random priorities. These measures were applied to both case studies, along with additional problem-specific measures developed from their formulations.

For BPPs, effective priorities reduced node counts by 52.7% to 58.8%, with improvements attributed to branching first on bin-usage variables and separating variables into classes. The random priorities provided no benefit. For IAM, no static ordering, whether rule-based or randomly permuted, beat Gurobi's default, and no instance property reliably predicted an effective ordering in advance. Network-based measures provided no significant benefit for VRPs. Notably, the BPP class-splitting method reduced both node counts and solve times by approximately 65% when applied to VRPs.

Simple structure-based priorities can substantially outperform default solving. However, given their effectiveness depends on formulation structure, they may not generalise to other instances of the same problem or to other problem types.

Autonomous Orchestration of Third-Party Optimisation Engines

20 Nov 9:20am-9:40am

Maz Lakadia
Evolve DS

Optimisation engines are widely used in airline crew planning to generate solutions under operational constraints and business priorities. In practice, decision-making rarely relies on a single optimisation run. Analysts execute multiple scenarios to explore trade-offs, assess parameter sensitivity, and justify outcomes for specific operational contexts.

This calibration and scenario management process is typically performed manually by specialised analysts, becoming increasingly time-consuming as operations grow in scale or complexity. Modern engines efficiently solve individual problems, but much of the surrounding exploratory work remains outside the optimiser.

This presentation describes an autonomous orchestration approach applied to a commercial third-party crew planning optimisation engine. The orchestration layer operates externally, treating the optimiser as a black box and interacting only through its standard interface. It autonomously decides which optimisation runs to execute, adjusts parameters, rules, and input data across scenarios within allowable bounds, and determines when sufficient solution quality has been achieved. The system automates multi-scenario sensitivity analysis and trade-off exploration in a single end-to-end workflow.

Although demonstrated for crew planning, the methods are optimiser-agnostic and grounded in established OR and meta-optimisation techniques. The approach produces structured, human-readable insights that reduce manual analyst effort and improve the speed and consistency of optimisation-driven decision-making.

Learning and Transferring Scheduling Heuristics through Genetic Programming for Dynamic Multi-Mode Projects

20 Nov 9:40am-10:00am

Yuan Tian, Yi Mei, Mengjie Zhang
Victoria University of Wellington

This talk presents a series of studies undertaken as part of a PhD thesis on genetic programming (GP) hyper-heuristics for dynamic multi-mode project scheduling, including an exploration of how evolved heuristic knowledge can guide large language models (LLMs).

Projects with limited resources must repeatedly decide which activities to start and how to execute them. When activities offer alternative execution modes and their durations are uncertain, these decisions must balance precedence requirements, resource competition and project completion time.

The research first develops GP representations for jointly selecting activities and execution modes. It then extends individual selections to resource-feasible activity groups, capturing interactions among concurrently executable activities. A knee-point filtering procedure reduces the computational cost of constructing candidate groups.

To improve learning efficiency, a surrogate-assisted framework characterises scheduling rules through their behaviour in representative decision situations. These characterisations support duplicate removal and offspring preselection before expensive simulation evaluations. Experiments show improved scheduling performance and, in later evolutionary stages,

approximately 40–55% fewer full fitness evaluations to reach comparable solution quality for most surrogate-assisted configurations.

Finally, the research investigates transferring GP-evolved knowledge to LLM-based scheduling through feature selection, feature hints and explicit rule guidance. Experiments on smaller projects show that selected forms of guidance improve scheduling quality over an unguided LLM, with different trade-offs between performance, token consumption and decision stability.

Together, these studies examine how evolutionary computation can learn effective scheduling heuristics, improve learning efficiency and provide domain knowledge for LLM-based decision making.

Recent developments in solving large linear programs

20 Nov
10:00am-
10:20am

Simon Bowly
Gurobi Optimization

Over the last few years, first-order methods, particularly primal-dual hybrid gradient (PDHG), have become a practical alternative to simplex and interior-point algorithms for linear programming. PDHG relies on matrix-vector products, making it well suited to GPUs, parallel computation, and very large problems. These methods complement established approaches, but typically produce lower-accuracy solutions and do not directly provide the basic solutions that are standard in the field.

This talk shares Gurobi's experience with first-order methods and explores how PDHG works alongside existing algorithms in practice. We combine PDHG with simplex and barrier in concurrent optimisation and use PDHG iterates to warm-start the search for higher-accuracy solutions. We discuss how these combinations work in practice and what they mean for users choosing a solution approach.

These combinations open up further opportunities for first-order methods and GPU computing, including potential use of PDHG within mixed-integer programming.

Session 7 – 10:40am-12:00pm

Chair: Andrea Raith

Who keeps the lights on? A two-stage stochastic model of field-workforce design for electricity network resilience in Auckland

20 Nov
10:40am-
11:00am

Ehsan Mahmoodi
University of Auckland

How long customers spend without power after a storm depends less on how many faults occur than on whether a crew can reach them. That is decided well in advance by where depots are, how many crews there are and with which skills, which shifts they work, and how much rest they need after a late call-out. We formulate field-workforce design for a distribution utility as a bi-objective two-stage stochastic program. The first stage chooses depots, crew numbers, skill mix, rosters, standby and storm-escalation rules. The second stage is the operational response to weather-driven faults, evaluated by simulation. The objectives are SAIDI and annual cost, and fatigue, backlog and off-rota call-outs are constrained. The model is calibrated on 1,625 real Vector fault interruptions from 2024/25 and ERA5 weather, and it predicts other years out of sample, including the 2023 storm year to within 1%. NSGA-II approximates the Pareto front. We derive scenario-wise lower bounds on SAIDI from information relaxations, including a time-indexed LP, and estimate them by sample average. The gap between simulated SAIDI and the bound measures the reliability a workforce could deliver but its operating rules do not. 24/7 rosters dominate, and the coupling between fatigue and resilience comes from the roster rather than crew numbers.

Continuous Representations of Pareto Fronts via High-Order KKT Expansions

20 Nov
11:00am-
11:20am

Jack Yarndley, Roberto Armellin
University of Auckland

Pareto fronts for nonlinear optimisation problems are generally approximated using discrete samples, each requiring the solution of an individual optimisation problem. We propose using differential algebra (DA) to expand the Karush-Kuhn-Tucker (KKT) optimality conditions with respect to trade-off parameters, producing high-order polynomial maps of locally optimal solutions and their objective values. These maps provide continuous approximations of local Pareto-front segments. These are then extended to the whole front utilising branch-aware automatic domain splitting (ADS), which manages estimated truncation error and changes in active constraints, while interval arithmetic provides bounds on polynomial quantities within each subdomain. This approach is applied to several astrodynamics problems, including direct transcriptions of low-thrust Earth–Mars transfers.

Column Generation Algorithm for Parallel Task Scheduling with Communication Delays

20 Nov
11:20am-
11:40am

Lewis Azzopardi, Emmanuel Araba, Oliver Sinnen
University of Auckland, Department of Electrical, Computer, and Software Engineering

Scheduling a task graph on identical parallel machines, where a communication delay applies whenever two dependent jobs run on different machines, is a central problem in parallel computing, and exact methods for

it have relied on state-space search and mixed-integer linear programming (MILP).

Delayed column generation, an optimisation technique for difficult MILP, has not been formulated for this model. Our work extends previous work from van den Akker et al. by integrating the consideration of communication delays through a same-machine term in the precedence constraints of the master problem.

Two solvers are built on it: a branch-and-price solver, which needs an exact assignment model where branching on job pairs stalls, and a hybrid search that uses column generation only for a lower bound and decides integer feasibility with that model. Both were compared with a traditional MILP solved by Gurobi on 6,978 combinations of 2,326 benchmark task graphs, with 4 to 500 jobs and two, four or eight machines. While the direct column generation method was not faster than the traditional MILP, the hybrid approach shows a much better scalability and allow for the first time to solve large problems with 100 or more jobs.

Experiments with Neighbourhood Pricing to speed up Column Generation

20 Nov
11:40am-
12:00pm

Andrew Mason, Basile Blayac, Andrea Raith
University of Auckland

Column generation is widely used to solve large-scale optimisation problems, but its practical effectiveness can be hindered by slow convergence. This paper investigates neighbourhood pricing as a means of accelerating convergence by restricting the feasible region of column-generation subproblems to neighbourhoods of existing columns. A neighbourhood-pricing column generation algorithm is developed for the Generalised Assignment Problem and evaluated on challenging benchmark instances from the literature. The results demonstrate substantial improvements over standard column generation, reducing both the number of iterations and the computational time per iteration without compromising solution quality.

Plenary – 1:00pm - 2:00pm

Chair: Andrew Mason

Say What, Not How: Constraint Modelling with MiniZinc

20 Nov
1:00pm-
2:00pm

Jip J. Dekker
Monash University

A standard approach to tackling operations research problems is to model them for a Mixed Integer Programming (MIP) solver. The real-world constraints must be carefully encoded as linear inequalities, introducing auxiliary variables and choosing formulations with strong linear relaxations. MiniZinc takes a different approach. It is a high-level, solver-independent modelling language in which you state the problem much as you would describe it: "at most three tasks can run simultaneously", "these assignments must all be different", "these routes must form a single tour". A compiler then translates the model for the solver of your choice, whether based on constraint programming (CP), Boolean satisfiability (SAT), local search, or MIP. That choice matters, but it is often difficult to know in advance which solver suits a problem best, and the answer can change as the model evolves or solvers improve. Keeping the high-level model independent of the solver lets you explore different solving approaches without reformulating the problem each time. This talk introduces MiniZinc through small examples and uses results from the annual MiniZinc Challenge to illustrate how solver strengths vary across problems. It closes with current work on incremental MiniZinc, supported by new incremental solver interfaces, allowing models to be updated and re-solved as requirements or data change without starting from scratch.

Session 8 – 2:20pm-3:00pm

Chair: David Robb

Spare Parts Inventory Management Utilizing Live Machine Health Data: A Digital Twin Approach

20 Nov
2:20pm-
2:40pm

Yale T. Herer, Jennifer Kruman
Technion - Israel Institute of Technology

We study a spare parts inventory system for multi-workstation manufacturing environments in which machine health deteriorates according to a continuous-time Markov chain and failures require replacement from a shared pool of spare parts. Building on Kruman et al. (2025), which assumed exponentially distributed replenishment lead times, we analyze the system under deterministic lead times—a setting that eliminates the memoryless property. To address this structural challenge, we develop the Marginal Benefit Ordering (MBO) heuristic, designed to be implemented within a digital twin that runs in parallel with the physical system. At each deterioration event, the digital twin simulates the system forward over the lead time horizon and orders one spare part if and only if the expected marginal benefit of having an additional part at the arrival time is positive. The

heuristic requires no preliminary calibration, e.g., of an optimal base-stock level, and thus can adapt automatically to changes in deterioration structure or system scale. We conduct an extensive computational test, across nine deterioration scenarios and three system sizes. We benchmark it against optimal static base-stock policies and ten known dynamic policies. MBO outperforms the optimal static base-stock policy by up to 24.4%, outperforms the best dynamic policy by up to 18.9%—without requiring exhaustive a priori simulations. We further show that MBO extends naturally to stochastic lead times where the heuristic continues to outperform the optimal static base-stock policy and the best dynamic policy in the large majority of configurations under exponential, Erlang, and lognormal lead-time distributions.

Benefit of the Doubt Models: Input Minimisation, Zero-Valued Measures and Weight Robustness

20 Nov
2:40pm-
3:00pm

Paul Rouse, Maryam Hasannasab, Dimitris Margaritis
University of Auckland

Benefit of the Doubt (BoD) models are special cases of Data Envelopment Analysis (DEA), traditionally formulated to maximise weighted combinations of outputs. Input-oriented BoD models provide an alternative formulation, particularly relevant when performance is assessed through resource minimisation rather than achievement maximisation. Hasannasab, Margaritis and Rouse (2026) developed output-oriented BoD models to identify the maximum common lower bound (α) on weights that individual decision-making units (DMUs) can sustain while preserving their unconstrained efficiency scores. They proposed using the lowest α among strongly efficient DMUs as a universal weight bound, preserving the efficiency frontier while revealing underperformance among less efficient DMUs. Building on this framework, the present study investigates whether a similar α-bound approach can be extended to input-oriented BoD models, with particular attention to the implications of zero-valued inputs for multiplier weights and the feasibility of positive α bounds. Preliminary findings indicate that zero-valued inputs can still receive positive multiplier weights, depending on the supporting hyperplanes associated with the efficiency frontier. We further investigate the conditions under which such positive weights arise and whether meaningful α bounds can be established when some input measures are zero.

Closing

Unuhia, unuhia	Draw on, draw on,
Unuhia ki te uru tapu nui	Draw on the supreme sacredness
Kia wātea, kia māmā, te ngākau, te tinana, te wairua i te ara takatā	To clear, to free the heart, the body and the spirit of mankind
Koia rā e Rongo, whakairia ake ki runga	Rongo, suspended high above us (i.e. in 'heaven')
Kia tina! TINA! Hui e! TāIKI E!	Draw together! Affirm!

Author Index

Abdelhadi
Ahmed, 10
Abdoli
Mohsen, 5
Araba
Emmanuel, 18
Armellin
Roberto, 17
Azzopardi
Lewis, 18

Blayac
Basile, 18
Bowly
Simon, 16

Costello
Seósamh B., 11

Dekker
Jip J., 19

Fargier
Alexandre, 8
Fonseca
Jose Da, 9

Gu
Jiaying, 11

Haas
Jannik, 7
Hasannasab
Maryam, 20
Herer
Yale T., 19
Hillas
Lisa Aoki, 5, 6

Ingimundardóttir
Helga, 6

Jaiswal
Samriddhi, 14

Kieu
Minh, 10
Kruman
Jennifer, 19
Kök
Ali, 8

Lakadia
Maz, 15
Lin
Amelia Fabia, 11
Liu
Hannah, 11

Mahmoodi
Ehsan, 17
Margaritis
Dimitris, 20
Marshall
Sarah, 5
Mason
Andrew, 18
Mei
Yi, 15

Newrick
Harry, 10

Oliver
Chris, 10

Peer
Rebecca A.M., 7
Philpott
Andy, 8, 13
Pisinger

David, 9
Purvis
Ava, 14

Raith
Andrea, 10, 12, 18
Rouse
Paul, 20
Rustad
Asta Susan, 9
Rúnarsson
Tómas P., 6

Satymov
Rasul, 7
Shi
Yajun, 9
Sigurjónsdóttir
Signý Kristín, 6
Sinnen
Oliver, 18
Steele
Joe, 6

Tian
Yuan, 15

Vaithianathan
Rhema, 6

Walker
Cameron, 6
Wang
Hanqing, 12
Yi, 11
Wotherspoon
Liam, 11

Yang
Yang, 12
Yao
Enjian, 12
Yarndley
Jack, 17

Zegeer
Madison L., 7
Zhang
Mengjie, 15
Tianyu, 12
Wenjun, 9