



ORSNZ ANNUAL CONFERENCE 2025

Operations Research Society of New Zealand (ORSNZ)

November 20-21, 2025

The University of Auckland, Auckland, New Zealand

Opening

Whakataka te hau ki te uru
Whakataka te hau ki te tonga

Kia mākinakina ki uta
Kia mātaratara ki tai

E hī ake ana te atakura

He tio, he huka, he hau hū

Tīhei mauri ora!

Cease the winds from the west
Cease the winds from the south

Let the breeze blow over the land
Let the breeze blow over the ocean

Let the red-tipped dawn come with a
sharpened air
A touch of frost, a promise of a glorious
day.

Welcome

Tēnā tatou e hoa ma

Greetings to us all

Kua haere mai nei

Who have come here

Ki te ako

To learn

I ngā whakaaro

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

Great is the thanks/appreciation

Mō tō koutou aronui

For your interest

Ki te aro mai

To attend

I ngā kaupapa

The programme

Otirā, kua mutu i konei ā ku mihi mō
tēnei wā

Thus, I end my greeting here for the
present.

Nō reira, tēnā koutou, tēnā koutou, ā,
tēnā tātou katoa.

Therefore, greetings and honour to one
and all.

Contents

Opening	i
Welcome	ii
Delegate Information	1
Programme Summary	2
Thursday 20th November 2025	4
Session 1 – 9:10am-10:10am	4
Don't Look Back in Anger: Wasserstein Distributionally Robust Optimization with Nonstationary Data (<i>Dominic Keehan, Edward Anderson, and Wolfram Wiesemann</i>)	4
Conex: A Computer Vision-Based System for Real-Time Monitoring of Temporary Traffic Management Setups (<i>Mostafa Papen, AT Computer Vision Team, and AT GIS team</i>)	5
Develop AI Models to Predict Water Levels (<i>Mahin Panchia, Elijah Hayward, Thomas Adams, and Cameron Walker</i>)	5
Session 2: 10:40am-11:40am	6
Gurobi on GPUs: Current State and Future Directions (<i>Hamideh Anjomshoa</i>)	6
Sequential Convex Programming for Multimode Spacecraft Trajectory Optimization (<i>Jack Yarnley</i>)	6
Session 3: Plenary – 11:30am-12:30pm	7
Algorithms for Constrained Optimization: The Benefits of General-purpose Software (<i>Michael Saunders</i>)	7
Session 4: Scheduling and assignment – 1:30pm-2:30pm	8
Accelerating Column Generation and devising Large Neighbourhood Search heuristics using Neighbourhood Pricing (<i>Basile Blayac, Andrew J. Mason, and Andrea Raith</i>)	8
Synergizing AI and Multi-Objective Optimization in Smart Manufacturing: Two Industrial Applications (<i>Ehsan Mahmoodi and Masood Fathi</i>)	8
HIOCS: Heuristic Inter-Operator Co-Scheduling Method for Efficient DNN Inference on GPUs (<i>Ning Wang, Andrea Raith, and Oliver Sinnen</i>)	9
Session 5: Sustainability and resilience– 2:40pm-3:40pm	10
Optimising Sustainable Meal Planning Under Real-World Constraints (<i>Juliette Foley</i>)	10

Modelling Energy Resilience in Tairāwhiti (Zainab Rizvi, Julia de Lange, Michael O’Sullivan, Natasha Koi, Harley Dibble, Holly Thorpe, and Cameron Walker)	10
Planning Sustainable Electric Taxi Fleets: From the Grid to the Grid (Ahmed Abdelhadi, Andrea Raith, and Minh Kieu)	11
Friday 21st November 2025	12
Session 6: Energy – 9:00am-10:20am	12
Exploring future electricity demand until 2075 in New Zealand using narrative scenarios (Catalina Klausen, Rafaella Canessa, Madison Zegeer, Akash Jyoti, Francisco Astorga, Meisy Fortunatus, Hadi Vatankhah, Vicente Sepulveda, Haolong Pang, Baxter Kamana, Stella Steidl, Karan Titus, Jannik Haas, and Rebecca Peer)	12
Integrating stochastic hydropower operational strategies into capacity expansion models: validation against SDDP (Madison L. Zegeer, Andy Philpott, Jannik Haas, and Rebecca A.M. Peer)	13
Multi-Period Aggregation of Distributed Energy Resources (Lisa Aoki Hillas and Zuguang Gao)	13
Optimizing grid-scale battery operations (Andy Philpott)	14
Session 7: Queueing – 10:50am-11:30am	14
Opaque service: the case of server selection by strategic customers (Binyamin Oz, Yoav Kerner, and Seva Shneer)	14
Determining the Optimal Queueing System Capacity in an Outpatient Center: a Queueing Approach (Mohsen Abdoli)	15
Session 8: Plenary – 11:40am - 12:40pm	15
Open-Source Transportation Modelling: From Research Prototype to Competitive Framework (Pedro Camargo)	15
Session 9: GPU computation – 1:30pm -2:30pm	16
MadSuite: GPU solvers for large-scale optimization (Alexis Montoison, François Pacaud, and Sungho Shin)	16
Recovering sparse DFT from missing signals via an interior point method on GPU (Vishwas Hebbur Venkata Subba Rao, Alexis Montoison, François Pacaud, Wei Kuang, and Mihai Anitescu)	16
One-shot optimisation: A framework for automated engine calibration and scenario management (Maz Lakadia)	17
Session 10: Supply chain and transport – 3:00pm -3:40pm	18
Supply Chain Company + More Optimal: Cloud-native Optimization for Faster, Transparent Supply Chains (Eddy de Haas and Siamak Moradi)	18
From Black Box to Insight: Diagnosing Unexpected Results of Airline Workforce Planning Models (Olga Perederieieva and Oliver Weide)	18
Closing	20
Author Index	21

Delegate Information

Conference Website

Detailed information about the conference is available on the conference website: <https://orsnz.org.nz/conf57/>

Venue

The conference will be held at the University of Auckland Conference Centre, in building 423, 22 Symonds Street. The sessions will be held in rooms 423-340 and 423-342.

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 6:30 pm on Thursday 20th November at the Fale Pasifika, University of Auckland <https://www.auckland.ac.nz/en/on-campus/life-on-campus/pacific-life/fale-pasifika.html>.

Sponsors

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



Waipapa
Taumata Rau
**University
of Auckland**



GUROBI
OPTIMIZATION



SUPPLY CHAIN
COMPANY



WePlan

ORSNZ Conference Programme

Thursday 20th November 2025

8:30 - 9:00 am	Registration
9:00 - 9:10 am	Opening (room 423-342)
9:10 - 10:10 am	Session 1 (room 423-342) (Chair: Cameron Walker) Dominic Keehan, University of Auckland, <i>Don't Look Back in Anger: Wasserstein Distributionally Robust Optimization with Nonstationary Data (YPP)</i> Mostafa Papen, Auckland Transport, <i>Conex: A Computer Vision-Based System for Real-Time Monitoring of Temporary Traffic Management Setups</i> Mahin Panchia and Elijah Hayward, University of Auckland, <i>Develop AI Models to Predict Water Levels (YPP)</i>
10:10 - 10:40am	Morning Tea
10:40 - 11:20 am	Session 2 (room 423-342) (Chair: Lisa Aoki Hillas) Hamideh Anjomshoa, Gurobi, <i>Gurobi on GPUs: Current State and Future Directions</i> Jack Yarnley, Te Pūnaha Ātea - Space Institute, University of Auckland, <i>Sequential Convex Programming for Multimode Spacecraft Trajectory Optimization (YPP)</i>
11:20 - 11:30 am	<i>Mini Break</i>
11:30 am - 12:30 pm	Session 3: Plenary (room 423-342) (Chair: Andy Philpott) Michael Saunders, Stanford University, <i>Algorithms for Constrained Optimization: The Benefits of General-purpose Software</i>
12:30 - 1:30 pm	<i>Lunch</i>
1:30 - 2:30 pm	Session 4: Scheduling and assignment (room 423-340) (Chair: Kevin Jia) Basile Blayac, University Of Auckland, <i>Accelerating Column Generation and devising Large Neighbourhood Search heuristics using Neighbourhood Pricing (YPP)</i> Ehsan Mahmoodi, University of Skövde, <i>Synergizing AI and Multi-Objective Optimization in Smart Manufacturing: Two Industrial Applications</i> Ning Wang, University of Auckland, <i>HIOCS: Heuristic Inter-Operator Co-Scheduling Method for Efficient DNN Inference on GPUs (JAG, YPP)</i>
2:30 - 2:40 pm	<i>Mini Break</i>
2:40 - 3:40 pm	Session 5: Sustainability and resilience (room 423-340) (Chair: Grant Read) Juliette Foley, University of Auckland, <i>Optimising Sustainable Meal Planning Under Real-World Constraints (JAG, YPP)</i> Zainab Rizvi, Julia de Lange, University of Auckland, <i>Modelling Energy Resilience in Tairāwhiti (JAG)</i> Ahmed Abdelhadi, University Of Auckland, <i>Planning Sustainable Electric Taxi Fleets: From the Grid to the Grid (JAG, YPP)</i>
3:40 - 4:00 pm	<i>Coffee break and walk to Fale Pasifika</i>
4:00pm	Gurobi Community Networking Event Fale Pasifika, University of Auckland (<i>requires separate registration</i>)
6:30 pm	Conference Dinner at Fale Pasifika, University of Auckland

Friday 6th December 2024

8:45 - 9:00 am	Registration
9:00 - 10:20 am	Session 6: Energy (room 423-340) (Chair: Andy Philpott) Catalina Klausen, University Of Canterbury, <i>Exploring future electricity demand until 2075 in New Zealand using narrative scenarios</i> (YPP) Madison L. Zegeer, University Of Canterbury, <i>Integrating stochastic hydropower operational strategies into capacity expansion models: validation against SDDP</i> (YPP) Lisa Aoki Hillas, University of Auckland, <i>Multi-Period Aggregation of Distributed Energy Resources</i> Andy Philpott, University Of Auckland, <i>Optimizing grid-scale battery operations</i>
10:20 - 10:50 am	<i>Morning Tea</i>
10:50 - 11:30 am	Session 7: Queueing (room 423-340) (Chair: Sarah Marshall) Binyamin Oz, University Of Auckland, <i>Opaque service: the case of server selection by strategic customers</i> (YPP) Mohsen Abdoli, University of Auckland, <i>Determining the Optimal Queueing System Capacity in an Outpatient Center: a Queueing Approach</i> (JAG, YPP)
11:30 - 11:40 am	<i>Mini Break</i>
11:40 am - 12:40 pm	Session 8: Plenary (room 423-340) (Chair: Andrea Raith) Pedro Camargo, Outer Loop Consulting and Argonne National Laboratory, <i>Open-Source Transportation Modelling: From Research Prototype to Competitive Framework</i>
12:40 - 1:00 pm	<i>Lunch</i>
1:00 - 1:30 pm	ORSNZ Annual General Meeting (room 423-340)
1:30 - 2:40 pm	Session 9: GPU Computation (room 423-340) (Chair: Andrew Mason) Alexis Montois, University of Chicago, <i>MadSuite: GPU solvers for large-scale optimization</i> Vishwas Hebbur Venkata Subba Rao, Argonne National Laboratory, <i>Recovering sparse DFT from missing signals via an interior point method on GPU</i> Maz Lakadia, Evolve DS Limited, <i>One-shot optimisation: A framework for automated engine calibration and scenario management</i>
2:30 - 3:00 pm	<i>Afternoon tea</i>
2:40 - 3:40 pm	Session 10: Supply chain and transport (room 423-340) (Chair: Mike O'Sullivan) Eddy de Haas, Supply Chain Company, <i>Supply Chain Company + More Optimal: Cloud-native Optimization for Faster, Transparent Supply Chains</i> Olga Perederieieva, WePlan, <i>From Black Box to Insight: Diagnosing Unexpected Results of Airline Workforce Planning Models</i>
3:40 - 4:00 pm	Prizegiving and Closing (room 423-340)

Thursday 20th November 2025

Session 1 – 9:10am-10:10am

Chair: Cameron Walker

Don't Look Back in Anger: Wasserstein Distributionally Robust Optimization with Nonstationary Data

Dominic Keehan, Edward Anderson, and Wolfram Wiesemann
University of Auckland

20 Nov
9:10am-
9:30am

We study data-driven decision-making problems where historical observations are generated by a time-evolving distribution whose consecutive shifts are bounded in Wasserstein distance. We address this nonstationarity using a distributionally robust optimization model with an ambiguity set that is a Wasserstein ball centered at a weighted empirical distribution, thereby allowing for the time decay of past data in a way which accounts for the drift of the data-generating distribution. Our main technical contribution is a concentration inequality for weighted empirical distributions that explicitly captures both the effective sample size (i.e., the equivalent number of equally weighted observations) and the distributional drift. Using our concentration inequality, we select observation weights that optimally balance the effective sample size against the extent of drift. The family of optimal weightings reveals an interplay between the order of the Wasserstein ambiguity ball and the time-decay profile of the weights. Classical weighting schemes, such as time windowing and exponential smoothing, emerge as special cases of our framework, for which we derive principled choices of the parameters. Numerical experiments demonstrate the effectiveness of the proposed approach.

Conex: A Computer Vision-Based System for Real-Time Monitoring of Temporary Traffic Management Setups

Mostafa Papen, AT Computer Vision Team, and AT GIS team
Auckland Transport

21 Nov
3:40pm-
4:00pm

Temporary traffic management (TTM) setups using traffic cones are crucial for ensuring road safety during construction or maintenance activities. However, manually monitoring these setups for compliance and detecting unauthorized installations remains a challenge. In this paper, we present Conex, a computer vision-based system that supports road safety by detecting TTM setups through traffic cone recognition and cross-referencing them with active notified TTM activities. The system uses camera snapshots to identify traffic cones, extract geolocation data, and validate configurations against authorized plans, addressing three key use cases: (1) identifying unauthorized TTM setups, (2) monitoring compliance with safety standards, and (3) detecting stray cones outside designated zones. The information generated by Conex can be shared with traffic management systems and auditors to assist in improving oversight and facilitating timely interventions. We discuss Conex’s implementation, including challenges such as occlusions and lighting variability, and evaluate its preliminary performance in real-world applications. Early findings suggest that systems like Conex could enhance monitoring of temporary traffic controls, reduce risks from non-compliant setups, and improve operational workflows. Further validation is required to assess its long-term impact on safety and efficiency.

Develop AI Models to Predict Water Levels

Mahin Panchia, Elijah Hayward, Thomas Adams, and Cameron Walker
University of Auckland

20 Nov
9:50am-
10:10am

Northland Regional Council (NRC) operates an extensive water-level monitoring network that is costly to maintain due to ongoing sensor calibration and servicing. This project investigated whether machine learning models could complement or partially replace physical sensors by predicting water levels using data from neighbouring sites and related environmental variables.

Data from Northland monitoring sites between 2020 and 2023 were obtained via the NRC Hilltop server, cleaned, and combined into a unified dataset. Models were developed starting with a multiple linear regression baseline and progressively introducing environmental features, spatial filtering, temporal lags, and nonlinear algorithms such as Gradient Boosting. Model performance was evaluated using R-squared, Mean Absolute Error, and Root Mean Square Error across multiple temporal splits.

Results showed the baseline linear model achieved strong accuracy at most sites, reflecting linear relationships between neighbouring locations. At more complex sites, incorporating rainfall, groundwater, spatial selection, and time-lagged predictors improved accuracy, while combining these with nonlinear algorithms highlighted the value of advanced modelling. Reliable accuracy generally required three to six months of training data, and periodic manual check-ins every three months further enhanced long-term predictions.

Overall, machine learning provides a practical, site-specific approach to enhancing NRC’s monitoring network. Linear regression offers an efficient solution for well-connected

sites, while nonlinear models suit complex environments. These models can also reconstruct missing data, improving reliability and reducing maintenance.

Session 2: 10:40am-11:40am

Chair: Lisa Aoki Hillas

Gurobi on GPUs: Current State and Future Directions

Hamideh Anjomshoa

Gurobi

20 Nov
10:40am-
11:00am

Recent research applying Primal-Dual Hybrid Gradient (PDHG) methods to linear programming has shown that first-order methods can be a useful part of the mathematical programming toolbox. These algorithms also lend themselves to GPU-based implementations, raising the prospect of exploiting the current surge in hardware development for optimisation applications.

This session will focus on the current state of Gurobi's GPU capabilities and explore future directions for leveraging GPUs in optimisation. Attendees will gain insights into how GPU technology can enhance the performance of Gurobi's solver and what advancements are on the horizon.

Sequential Convex Programming for Multimode Spacecraft Trajectory Optimization

Jack Yarnley

Te Pūnaha Ātea - Space Institute, University of Auckland

20 Nov
11:00am-
11:20am

Spacecraft equipped with multiple propulsion modes or systems can offer enhanced performance and mission flexibility compared with traditional configurations. Each mode provides distinct trade-offs between thrust level and efficiency, or may correspond to entirely different propulsion technologies. This paper presents a sequential convex programming (SCP) approach for the optimal design of multi-mode and multi-propulsion spacecraft trajectories. The method extends the dynamical linearization within SCP using sparse automatic differentiation, enabling efficient inclusion of multiple propulsion modes without complex manual reformulation while maintaining comparable computational efficiency. Several new lossless constraint formulations are introduced to ensure selection of a single propulsion mode at each time step and limit the total number of modes used. The approach is demonstrated for (1) a low-thrust Earth-67P rendezvous using the SPT-140 thruster with 21 discrete modes, and (2) an Earth-Mars transfer employing both a low-thrust engine and a solar sail. Results confirm that the proposed method can efficiently compute optimal trajectories for these scenarios.

Session 3: Plenary – 11:30am-12:30pm

Chair: Andy Philpott

Algorithms for Constrained Optimization: The Benefits of General-purpose Software

Michael Saunders

Stanford University

20 Nov
11:30am-
12:30pm

We review the history of numerical optimization at Stanford University and describe some unexpected applications of optimization software within aerospace, signal analysis, systems biology, economics, radiotherapy,

Sometimes general-purpose software leads to new applications. Sometimes new applications lead to new algorithms (which we implement with general-purpose software).

Session 4: Scheduling and assignment – 1:30pm-2:30pm

Chair: Kevin Jia

Accelerating Column Generation and devising Large Neighbourhood Search heuristics using Neighbourhood Pricing

Basile Blayac, Andrew J. Mason, and Andrea Raith
University Of Auckland

20 Nov
1:30pm-
1:50pm

In the context of column generation, Neighbourhood Pricing is a column pricing strategy that restricts the search for new columns to the neighbourhoods of existing ones. Our study focuses on the Generalized Assignment Problem (GAP), a well-known and challenging scheduling problem where tasks must be assigned to heterogeneous agents under capacity constraints. Two variants of Neighbourhood Pricing for the GAP are proposed, along with methods to use them effectively. We demonstrate the stabilisation effects of Neighbourhood Pricing to solve root nodes of literature instances. Our study shows that this approach exploits structural properties of literature instances, and how these properties can be leveraged further. Finally, we study neighbourhoods of integer solutions defined using Neighbourhood Pricing and compare them with other types of neighbourhoods commonly used in matheuristic approaches.

Synergizing AI and Multi-Objective Optimization in Smart Manufacturing: Two Industrial Applications

Ehsan Mahmoodi and Masood Fathi
University of Skövde

20 Nov
1:50pm-
2:10pm

Modern manufacturing systems require advanced decision support frameworks that integrate the analytical and predictive capabilities of artificial intelligence (AI) with the prescriptive strengths of multi-objective optimization. This study introduces two complementary applications that exemplify this integration within operational manufacturing settings. The first application utilizes data-driven simulation in combination with the Non-dominated Sorting Genetic Algorithm II (NSGA-II) for resource allocation in a marine engine manufacturer in Sweden. In this context, AI-based clustering algorithms are used to conduct post-optimality analysis of the solution space, which identifies key relationships between decision variables and system throughput. The second application focuses on buffer and resource allocation in high-mix, low-volume hybrid flow shops. It employs machine learning regressors as surrogate models to reduce the computational cost of simulation evaluations. In this scenario, AI supports lifelong meta-learning for adaptive optimization across different problem instances, enables intelligent operator selection, and applies pattern-mining algorithms to establish connections between solution configurations and objective performance. Together, these applications illustrate the potential of AI to enhance multi-objective optimization at various stages, including solution evaluation, search guidance, and knowledge extraction from Pareto-optimal solutions.

HIOCS: Heuristic Inter-Operator Co-Scheduling Method for Efficient DNN Inference on GPUs

Ning Wang, Andrea Raith, and Oliver Sinn
University of Auckland

20 Nov
2:10pm-
2:30pm

Deep neural networks (DNNs) are increasingly deployed in real-time applications, yet their inference performance is often constrained by inefficient GPU utilization. While various inter-operator scheduling methods have been proposed to address this issue, many suffer from coarse-grained operator classifications, incomplete parallelism, unawareness of GPU resource contention, etc. To overcome these limitations, we propose HIOCS, a heuristic inter-operator co-scheduling framework that accelerates DNN inference on GPUs through fine-grained kernel-level scheduling. HIOCS employs a graph optimization technique at kernel granularity that reconstructs data dependencies based on memory access patterns, thereby uncovering latent parallelism. Building on this optimized graph, a heuristic co-scheduling algorithm is introduced to prioritize critical-path kernels and selectively co-locate non-critical kernels by jointly considering latency impact, resource contention, and kernel affinity to minimize makespan and maximize GPU utilization. Extensive experiments on popular DNN models across diverse domains demonstrate that HIOCS consistently achieves higher GPU utilization and better scalability. Compared to PyTorch default mode and the state-of-the-art method Opara, HIOCS achieves up to 21.19 $\times$ and 1.13 $\times$ speedup in inference latency, respectively, making it well-suited for latency-sensitive deployment scenarios.

Session 5: Sustainability and resilience— 2:40pm-3:40pm

Chair: Grant Read

Optimising Sustainable Meal Planning Under Real-World Constraints

Juliette Foley

University of Auckland

20 Nov
3:00pm-
3:20pm

Food systems account for 26% of global greenhouse gas emissions, yet consumers lack practical tools for sustainable meal planning. While operations research has addressed diet optimisation through linear programming, and recent applications provide carbon footprint tracking, existing approaches fail to integrate heterogeneous recipe data, temporal price volatility, and discrete purchasing constraints (e.g., one cannot purchase 1 tablespoon of milk).

This work presents an end-to-end computational framework integrating: (1) automated recipe processing with statistical weight extraction from ingredient text; (2) graph-based emissions database mapping; (3) hybrid fuzzy-LLM ingredient-to-product matching; (4) temporal price modelling across 37 weeks and 8,000+ products; (5) bi-objective mixed-integer linear programming that selects optimal recipe-ingredient-product combinations. With a single optimisation solution corresponding to a feasible shopping list and a recipe set, which recipes are chosen, which retail products are used for each ingredient, and the purchase quantities required to satisfy the recipes.

Applied to 6,692 recipe-ingredient pairs, the system achieves 99.8% ingredient matching, 95.2% emissions coverage, and 94% solution feasibility under real-world price volatility. A robustness analysis spanning 37 weeks reveals that optimal recipe selection adapts to promotional pricing. Representative Pareto frontier analysis demonstrates that 30-40% emissions reductions are achievable at 12-18% cost increases, with three distinct trade-off regions exhibiting cost-emission gradients of \$3.80–\$12.40/kg CO₂e depending on nutritional constraints.

Overall, this work bridges a key research-to-practice gap, helping to facilitate sustainable meal planning within household budget constraints.

Modelling Energy Resilience in Tairāwhiti

Zainab Rizvi, Julia de Lange, Michael O’Sullivan, Natasha Koi, Harley Dibble, Holly Thorpe, and Cameron Walker

University of Auckland

20 Nov
3:20pm-
3:40pm

Energy hardship and resilience are key concerns in Tairāwhiti, both because of and despite climate change. Distributed energy resources have the potential to provide resilience and the opportunity to alleviate said hardship for households. In this work, we present the transdisciplinary research we undertook as we built partnerships with people from communities in Tairāwhiti. We present the evolution of this partnership as we identify the energy trilemma challenges of affordability, sustainability, and reliability faced by communities in Tairāwhiti, along with the modelling that empowers our partners to explore solutions.

In dialogue with our community partners, we used operations research models to: estimate the effect on national generation and transmission of various levels of uptake of local generation; and simulate how local renewable generation, battery storage, and energy sharing arrangements can enhance energy security and affordability.

Furthermore, we report on communities building resilience into their energy system through local renewable energy generation. Our models provide a virtual space to test innovative energy sharing strategies from the increased generation, and allow us to observe their implications. Lastly, we present further development in our transdisciplinary research through the ensuing questions being asked by our community partners.

Planning Sustainable Electric Taxi Fleets: From the Grid to the Grid

Ahmed Abdelhadi, Andrea Raith, and Minh Kieu
University Of Auckland

20 Nov
2:40pm-
3:00pm

This paper investigates the operational and infrastructure implications of integrating vehicle-to-grid (V2G) technology into a wireless charging network for an electric taxi fleet in Auckland, New Zealand. We develop a discrete-event simulation combined with a mixed-integer linear programming model to jointly optimise charging station placement, pile allocation, and depot upgrades under realistic demand profiles. The framework explicitly accounts for battery degradation using a semi-empirical Li-NMC model incorporating both cyclic and calendar ageing. Two scenarios are evaluated over a three-year horizon where out-of-service charging is restricted to low grid load periods only: V2G disabled and V2G enabled with 20 kWh AC static wireless charging (SWC) infrastructure. Results show that V2G participation increases total grid-to-vehicle energy delivery by 14%, however 9-11% of monthly energy was returned back to the grid during high grid load periods. Opportunistic SWC charging demonstrated its potential to mitigate range anxiety and support continuous operations even with V2G participation. However, V2G introduces modest service trade-offs: fulfilled trips dropped by 2.3%, and average passenger waiting time rises from 11 to 14 minutes. However, battery health remains robust in both cases, with state of health above 98% after 36 months, confirming minimal additional degradation from V2G cycling. These findings suggest that V2G is technically feasible for commercial fleets.

Friday 21st November 2025

Session 6: Energy – 9:00am-10:20am

Chair: Andy Philpott

Exploring future electricity demand until 2075 in New Zealand using narrative scenarios

21 Nov
9:00am-
9:20am

Catalina Klausen, Rafaella Canessa, Madison Zegeer, Akash Jyoti, Francisco Astorga, Meisy Fortunatus, Hadi Vatankhah, Vicente Sepulveda, Haolong Pang, Baxter Kamana, Stella Steidl, Karan Titus, Jannik Haas, and Rebecca Peer

University Of Canterbury

New Zealand has committed to net-zero greenhouse gas emissions by 2050. Emerging energy-intensive users, such as data centres and CO₂ capture, or electrifiable high-emitting sectors, such as transport and heating, could compromise New Zealand's targets. To plan the energy transition, operations research approaches are often employed, especially optimisation on a least-cost basis. These modelling efforts are directly informed by future projections of energy demand. However, previous national studies have not included potential new demand drivers in their future demand projections, potentially underestimating future demand.

To address this gap, this work aims to quantify different demand scenarios for future energy planning that could enable New Zealand to reach net-zero emissions by 2050. This study presents four scenarios: Business as Usual, Delayed Transition, Accelerated with Local Innovation, and Cheap Electricity, each representing possible future energy demand, based on technology adoption curves, efficiency improvements, and emission targets. Using these parameters, this study projects the energy demand for eight different energy drivers until 2075, including transportation, residential demand, industrial heating, chemical feedstocks, desalination, emerging proteins, data centres, and carbon capture and sequestration. Results show the range of possible energy demand futures for New Zealand under our four narrative scenarios. Carbon capture and sequestration emerges as a potentially particularly influential energy demand driver to enable a net-zero emissions economy.

Integrating stochastic hydropower operational strategies into capacity expansion models: validation against SDDP

Madison L. Zegeer, Andy Philpott, Jannik Haas, and Rebecca A.M. Peer
University Of Canterbury

21 Nov
9:20am-
9:40am

Hydro-dominated energy systems provide valuable operational flexibility through their combined generation and storage capacity, yet they are increasingly exposed to hydrological variability intensified by climate change. In New Zealand, thermal generation has historically buffered this variability, but as fossil fuels are phased out, new strategies are needed to manage uncertainty. The challenge lies in linking detailed hydropower operations, well captured by high-fidelity scheduling models, with long-term capacity expansion planning, where operational realism is difficult to incorporate efficiently.

This study validates the GEMSTONE capacity expansion model against JADE, a Stochastic Dual Dynamic Programming (SDDP) hydrothermal scheduling model. GEMSTONE employs a reduced-order stochastic formulation that embeds inflow uncertainty through simplified reservoir release policies, allowing tractable integration into long-term planning frameworks. JADE serves as a high-fidelity benchmark, representing dynamic reservoir behaviour and short-term hydrological variability.

Model calibration across multiple inflow scenarios, demand profiles, and system capacities enables direct comparison of dispatch and cost outcomes. Results show that GEMSTONE reproduces JADE's hydro dispatch and system operating costs with strong agreement, despite its coarser temporal and stochastic resolution. Differences are observed in storage dynamics: GEMSTONE maintains higher average reservoir levels due to its uniform release strategy, whereas JADE adapts more flexibly to wet and dry conditions.

Overall, this validation demonstrates that reduced-order stochastic formulations can approximate high-fidelity hydrothermal scheduling models. GEMSTONE provides a tractable, operations research-based framework for incorporating hydrological uncertainty into long-term capacity expansion planning. Future work includes extending this framework to include climate-driven inflow shifts and with additional long-duration energy storage options.

Multi-Period Aggregation of Distributed Energy Resources

Lisa Aoki Hillas and Zuguang Gao
University of Auckland

21 Nov
9:40am-
10:00am

This paper looks at the problem of how distributed energy resources (DERs), such as household-level solar panels, should be integrated into the wholesale electricity market. We assume that DER owners have access to batteries that allow them to transfer power across time periods. We look at how DER owners should choose to consume, exchange with the market, and buy or sell power. Additionally, given the choices of the DER owners, we look at how an aggregator whose role is to procure electricity from DERs and sell the electricity in the wholesale market should set prices to the DER owners.

Optimizing grid-scale battery operations

Andy Philpott

University Of Auckland

21 Nov
10:00am-
10:20am

Large battery energy storage systems (BESS) are being built and commissioned in the New Zealand electricity market. Recently commissioned examples are the Rotowaro 35 MWh BESS at Huntly and Meridian's 200 MWh Ruakaka BESS in Northland. Their value accrues from a number of different uses, e.g. shifting load, or providing reserve, or supporting frequency and voltage. We focus on their ability to earn revenue from shifting load from high priced periods to low priced periods. Batteries can operate behind the meter, where they shift load by arbitraging wholesale electricity prices, or they can offer supply curves (specifying buy and sell prices) to SPD, the electricity market pricing and dispatch model. We contend that the second approach produces more efficient outcomes, but there are reasons why battery operators might prefer the first approach (even though time-averaged prices for each half hour makes it difficult to arbitrage within this period). We argue for a change to SPD that will incentivize batteries to offer their load shifting capacity as state dependent supply curves.

Session 7: Queueing – 10:50am-11:30am

Chair: Sarah Marshall

Opaque service: the case of server selection by strategic customers

Binyamin Oz, Yoav Kerner, and Seva Shneer

University Of Auckland

21 Nov
10:50am-
11:10am

We analyze a multi-server strategic queueing model in which customers strategically choose among servers with differing service rates and quality valuations. Customers can either commit to a specific server or adopt a flexible strategy, simultaneously joining multiple queues and ultimately receiving service from the server that first completes their request, after which all remaining redundant requests are canceled. This flexibility option presents a clear trade-off: customers benefit from potentially shorter waiting times, but risk receiving lower-quality service from less desirable servers. Additionally, the effectiveness of flexibility depends intricately on the decisions of customers who commit to specific servers, thereby complicating strategic interactions.

We characterize the Nash equilibrium outcomes of this strategic queueing game, highlighting how customers balance waiting time reduction against the uncertainty in service quality. Crucially, our analysis demonstrates that incorporating the flexibility option leads to higher overall social welfare, even when customers individually act to maximize their own utilities. We further examine the socially optimal allocation and the revenue maximizing pricing strategy, when flexibility is offered at a price.

Determining the Optimal Queuing System Capacity in an Outpatient Center: a Queuing Approach

Mohsen Abdoli
University of Auckland

21 Nov
11:10am-
11:30am

Determining the optimal queuing system capacity is a critical challenge in outpatient clinics, as it directly affects patient waiting time, patient rejection, and physician idle-time costs. This study employs an M/M/C/K queuing model to identify the system capacity that minimizes the expected total cost. To validate the queuing model, a discrete event simulation of the clinic was developed, and its results were compared with those of the analytical model. The analysis considers two scenarios reflecting different patient behaviors. The first scenario represents current clinic conditions, where patients exhibit low sensitivity to waiting due to low fees. In this case, the objective function includes the expected costs of patient waiting and patient rejection. The second scenario assumes that patients may become more sensitive to delays, influenced by changes in insurance policies or the availability of public clinics with higher service quality. Accordingly, the expected cost of physician idle-time is incorporated into the objective function alongside the other costs. Optimal solutions for both scenarios were obtained using the full counting method. Subsequently, the study examines how each cost component varies with system capacity and performs sensitivity analyses on key parameters affecting the model. Finally, the results of the two scenarios are compared to provide deeper managerial insights, enabling clinic administrators to balance patient satisfaction, service efficiency, and operational costs. This approach offers a systematic framework for improving outpatient clinic performance and guiding decision-making in capacity planning.

Session 8: Plenary – 11:40am - 12:40pm

Chair: Andrea Raith

Open-Source Transportation Modelling: From Research Prototype to Competitive Framework

Pedro Camargo
Outer Loop Consulting and Argonne National Laboratory

21 Nov
11:40am-
12:40pm

Transport modelling is a multi-disciplinary field focused on understanding human behaviour and their interactions with the built environment. Despite being over 70 years old and used all over the world, commercial software providers have enjoyed no competition from Open-Source until very recently. Born from the need for research tools, AequilibraE has evolved into a competitive alternative to commercial software, despite the modest investments it receives. The quickly growing adoption of the software by research institutions, government entities and private companies validates the need for an Open-Source solution in the field and the design decisions made so far.

From high-performance algorithmic solutions to carefully crafted documentation and concerns with DevOps, this talk will retrace the steps that ensured AequilibraE's success and the most important algorithmic components behind the software.

Session 9: GPU computation – 1:30pm -2:30pm

Chair: Andrew Mason

MadSuite: GPU solvers for large-scale optimization

Alexis Montoison, François Pacaud, and Sungho Shin
University of Chicago

21 Nov
1:30pm-
1:50pm

In recent years, the development of scalable continuous optimization solvers on GPUs has made significant progress, primarily driven by advances in GPU-based sparse linear algebra. This talk presents a suite of Julia-based open-source solvers designed to efficiently tackle large-scale linear, quadratic, and nonlinear programming problems with full GPU acceleration.

We first introduce MadNLP.jl, a GPU-native solver for nonlinear programming (NLP) based on interior-point methods and sparse direct solvers. It was the first solver in the suite and remains central to solving general nonlinear problems efficiently. Building on this foundation, MadNCL.jl is a robust meta-solver that orchestrates multiple NLP solves to handle degeneracy, ill-conditioning, and to achieve high-accuracy solutions with tolerances below $1e-8$. The latest solver MadIPM.jl, targets large-scale linear and convex quadratic programs (LP / QP). All three solvers are based on second-order algorithms, enabling greater robustness and accuracy compared to purely first-order methods such as PDLF. We present performance results on real-world benchmark instances, demonstrating substantial speedups between CPU and GPU implementations, while maintaining high solution quality.

Recovering sparse DFT from missing signals via an interior point method on GPU

Vishwas Hebbur Venkata Subba Rao, Alexis Montoison, François Pacaud, Wei Kuang, and Mihai Anitescu
Argonne National Laboratory

21 Nov
1:50pm-
2:10pm

We propose a method to recover the sparse discrete Fourier transform (DFT) of a signal that is both noisy and potentially incomplete, with missing values. The problem is formulated as a penalized least-squares minimization based on the inverse discrete Fourier transform (IDFT) with an ℓ_1 -penalty term, reformulated to be solvable using a primal-dual interior point method (IPM). Although Krylov methods are not typically used to solve Karush-Kuhn-Tucker (KKT) systems arising in IPMs due to their ill-conditioning, we employ a tailored preconditioner and establish new asymptotic bounds on the condition number of preconditioned KKT matrices. Thanks to this dedicated preconditioner – and the fact that FFT and IFFT operate as linear operators without requiring explicit matrix materialization – KKT systems can be solved efficiently at large scales in a matrix-free manner. Numerical results from a Julia implementation leveraging GPU-accelerated interior point methods, Krylov methods, and FFT toolkits demonstrate the scalability of our approach on problems with hundreds of millions of variables, inclusive of real data obtained from the diffuse scattering from a slightly disordered Molybdenum Vanadium Dioxide crystal.

One-shot optimisation: A framework for automated engine calibration and scenario management

Maz Lakadia
Evolve DS Limited

21 Nov
2:10pm-
2:30pm

Optimisation engines have been used for decades in flight operations to automatically generate solutions following many different rules and priorities. While this work is done very well by most engines, there's a lot more to the job than generating one solution with one set of rules and priorities.

It's common to generate multiple solutions with different trade-offs to refine a solution for a particular instance, or to get insights on rule / parameter sensitivity. This requires extra engine calibration and scenario management (creation and comparison). Generally, this work is done by a team of analysts who need to understand both OR and flight operations.

In this presentation we propose a framework for automating a series of optimisations (including automation of parameter adjustments) to get the best solution, along with corresponding insights / justifications in "one shot". Whilst the motivation for this work comes from flight operations, the general framework is applicable to many operational usages of optimisation engines.

Session 10: Supply chain and transport – 3:00pm - 3:40pm

Chair: Mike O’Sullivan

Supply Chain Company + More Optimal: Cloud-native Optimization for Faster, Transparent Supply Chains

Eddy de Haas and Siamak Moradi

Supply Chain Company

21 Nov
3:00pm-
3:20pm

Supply Chain Company (SCC) and More Optimal have combined domain expertise and a cloud-native optimization platform to deliver a powerful decision-support stack for supply-chain teams. The integrated offering pairs SCC’s track record deploying RELEX and Optimity across food, FMCG and manufacturing networks with More Optimal’s low-code modeling layer, solver-backed execution and interactive visualizations.

Key capabilities include: a low-code surface to encode business rules and custom constraints; a scalable, parallelized solver engine for routing, inventory, scheduling and batching; APIs and data adapters for reproducible scenario sampling and streaming forecasts; and built-in reporting for trade-off analysis and optimality-gap transparency. The platform supports rapid prototyping of MILP/CP models, hybrid heuristic-exact workflows, and production rollouts with automated monitoring and versioned deployments.

In production benchmarks, the stack reduces total cost and inventory days while cutting solve time versus monolithic approaches, without sacrificing transparency or control over risk and sustainability metrics. Teams benefit from faster iteration, clearer decision tradeoffs and easier handoffs from pilot to operations.

We’re inviting supply-chain modelers, data teams and deployment leads to pilot the integrated solution on real datasets — to accelerate model development, scale to industrial instance sizes, and unlock measurable operational improvements.

From Black Box to Insight: Diagnosing Unexpected Results of Airline Workforce Planning Models

Olga Perederieieva and Oliver Weide

WePlan

21 Nov
3:20pm-
3:40pm

Airline workforce planning must satisfy flight coverage, qualifications, training, and strict duty-time rules at scale. Optimization is effective, but large mixed-integer models often act as black boxes: when a solve is infeasible or produces unexpected assignments or costs, the root cause is hard to find.

We present methods in WePlan airline workforce planning optimizer to explain infeasibility and surprising outcomes. First, we use irreducible infeasible subset (IIS) analysis to isolate minimal conflicting constraints and link them to airline concepts (flight coverage, legal duty/rest limits, base and qualification rules, training windows). Although tools such as Gurobi’s IIS provide valuable signals, they are not always sufficient for end users due to granularity, size, or interpretability. Second, we run structured re-solves of modified

models: selectively relaxing constraints, adding penalized slacks, or toggling assumptions (e.g., training timing, reserve levels, base capacity). Comparing outcomes highlights which rules drive conflicts and what small changes restore feasibility.

We will present results from our initial investigations applying IIS-based diagnostics and structured re-solves in airline crew planning, demonstrating improved transparency, faster analysis, and greater confidence in optimization.

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, 11
Abdoli
 Mohsen, 15
Adams
 Thomas, 5
Anderson
 Edward, 4
Anitescu
 Mihai , 16
Anjomshoa
 Hamideh, 6
Astorga
 Francisco, 12

Blayac
 Basile, 8

Camargo
 Pedro, 15
Canessa
 Rafaella, 12

de Haas
 Eddy, 18
de Lange
 Julia, 10
Dibble
 Harley, 10

Fathi
 Masood , 8
Foley
 Juliette, 10
Fortunatus
 Meisy, 12

Gao
 Zuguang, 13

Haas
 Jannik, 12, 13
Hayward
 Elijah, 5
Hebbur Venkata Subba Rao
 Vishwas, 16
Hillas
 Lisa Aoki, 13

Jyoti
 Akash, 12

Kamana
 Baxter, 12
Keehan
 Dominic, 4
Kerner
 Yoav, 14
Kieu
 Minh, 11
Klausen
 Catalina, 12
Koi
 Natasha, 10
Kuang
 Wei, 16

Lakadia
 Maz, 17

Mahmoodi
 Ehsan, 8
Mason
 Andrew, 8
Montoison
 Alexis, 16
Moradi
 Siamak, 18

O'Sullivan	Sungho, 16
Michael, 10	Shneer
Oz	Seva, 14
Binyamin, 14	Sinnen
Pacaud	Oliver, 9
François, 16	Steidl
Panchia	Stella, 12
Mahin, 5	Thorpe
Pang	Holly, 10
Haolong, 12	Titus
Papen	Karan, 12
Mostafa, 5	Vatankhah
Peer	Hadi, 12
Rebecca, 12, 13	Walker
Perederieieva	Cameron, 5, 10
Olga, 18	Wang
Philpott	Ning, 9
Andy, 13, 14	Weide
Raith	Oliver , 18
Andrea, 8, 9, 11	Wiesemann
Rizvi	Wolfram, 4
Zainab, 10	Yarndley
Saunders	Jack, 6
Michael, 7	Zegeer
Sepulveda	Madison L., 12, 13
Vicente, 12	
Shin	