Abstracts of the PhD presentations

Ruben van Beesten
University of Groningen
e.r.van.beesten@rug.nl
Supervisors Ward Romeijnders and Kees Jan Roodbergen

Title *Solving mixed-integer risk-averse stochastic programs using convex approximations*
Abstract Many real-life optimization problems are stochastic in nature. In such uncertain environments, not only the *expected* outcome, but also the level of *risk* associated with a decision are often important criteria for determining the best decision. We study a class of optimization models that explicitly incorporate this risk-aversion in the objective: *two-stage risk-averse stochastic programs.* Here, the first stage consists of making a decision under only distributional knowledge of several random parameters. In the second stage, after observing the realization of the random parameters, we can make corrective actions in order to satisfy all constraints.
We focus on problems with integer restrictions in the second-stage problem. Such restrictions arise naturally in many practical settings (e.g., indivisibilities, fixed batch sizes, etc.). The resulting *mixed-integer* second-stage problems are generally non-convex. As a result, our optimization problem is generally not convex and hence, very hard to solve.
We propose to overcome this issue by means of a *convex approximation* of the original non-convex optimization model. Solving this convex approximation model is much easier, since we can make use of the rich toolbox of convex optimization. This provides us with an approximate solution to the original model. Clearly, this approach is not exact and we will inevitably make an approximation error. Therefore, we derive an upper bound on this approximation error, which serves as a performance guarantee for our solution method.

Mark van den Bergh
Leiden University
m.j.h.van.den.bergh@math.leidenuniv.nl
Supervisors Floske Spieksma and Walter Kosters

Title *Synchronized Hackenbush games*
Abstract The field of Combinatorial Game Theory studies two-player, deterministic, perfect information games, providing a beautiful framework in which these games can be analyzed. We wonder what happens to this framework if one of the three assumptions is dropped -- in particular, the assumption of perfect information. In this talk, we will consider the game of Hackenbush, in which players take turns removing edges from a colored graph; the last player to move, wins. To introduce imperfect information, we let the players move simultaneously, instead of taking turns. By doing so, optimal strategies are no longer necessarily deterministic, and we need to resolve to finding Nash equilibria to analyze the game. However, we will show that, in some sense, this 'synchronized' version of Hackenbush still closely resembles the original game.

Jasper Bos
University of Twente
j.bos@utwente.nl
Supervisors Richard Boucherie and Maartje van de Vrugt

Title *The Vehicle Routing Problem with Urgent Stochastic Customers*

Abstract Several industries deal with a combination of routine, plannable tasks and incident response. Examples are parcel pickup and delivery, maintenance, surveillance and security. Our research is inspired on hospital cleaning. In the Jeroen Bosch Ziekenhuis, routine tasks and incident response are divided over two teams. As travel distance as well as the response time are important performance measures, intuitively it might be beneficial to distribute the cleaners over routine tasks such that the expected response time to incidents is minimized. With data of previous incident times and locations, we can further improve the distribution of cleaners with respect to incidents. The main research questions are how to route and schedule the routine tasks while optimizing for both travel time and response time and whether it is beneficial to merge the teams or not.
We have developed a mixed integer program that solves the A priori Vehicle Routing Problem with Urgent Stochastic Customers. Ongoing work involves a Genetic Algorithm to solve instances of realistic size. Additionally, we simulated different scenarios and model parameters to investigate the trade-off between travel time and incident response time.

Ruben Brokkelkamp
CWI
Ruben.brokkelkamp@cwi.nl
Supervisor Guido Schäfer

Title *Approximate Pricing in Networks*
Abstract We introduce and study two new pricing problems in networks: Suppose we are given a directed graph $G = (V, E)$ with non-negative edge costs$(c_e)_{e\in E}$, k commodities $(s_i, t_i, w_i)_{i\in[k]}$ and a designated node $u \in V$. Each commodity $i \in [k]$ is represented by a source-target pair $(s_i, t_i) \in V \times V$ and a demand $w_i > 0$, specifying that w_i units of flow are sent from s_i to t_i along shortest s_i, t_i-paths (with respect to $(c_e)_{e\in E}$). The demand of each commodity is split evenly over all shortest paths. Assume we can change the edge costs of some of the outgoing edges of u, while the costs of all other edges remain fixed; we also say that we *price* (or *tax*) the edges of u.

We study the problem of pricing the edges of u with respect to the following two natural objectives: (i) *max-flow*: maximize the total flow passing through u, and (ii) *max-revenue*: maximize the total revenue (flow times tax) through u. Both variants have various applications in practice. For example, the max flow objective is equivalent to maximizing the *betweenness centrality* of u, which is one of the most popular measures for the influence of a node in a (social) network. We prove that (except for some special cases) both problems are NP-hard and inapproximable in general and therefore resort to approximation algorithms. We derive approximation algorithms for both variants and show that the derived approximation guarantees are best possible.

Melvin Drent
University of Luxembourg
melvin.drent@uni.lu
Supervisor Joachim Arts

Title *Expediting in Two-Echelon Spare Parts Inventory Systems*
Abstract We consider a two-echelon spare parts inventory system consisting of one central warehouse and multiple local warehouses. Each warehouse keeps multiple types of repairable parts to maintain several types of capital goods. The local warehouses face Poisson demand and are replenished by the central warehouse. We assume that demands that cannot be satisfied from stock at the central warehouse are back-ordered. The repair shop at the central warehouse has two repair options for each repairable part: a regular repair option and an expedited repair option. Irrespective of the repair

option, each repairable part uses a certain resource for its repair. Assuming a dual-index policy at the central warehouse and base stock control at the local warehouses, an exact and efficient evaluation procedure for a given control policy is formulated. To find an optimal control policy, we look at the minimization of total investment costs under constraints on both the aggregate mean number of backorders per capital good type and the aggregate mean fraction of repairs that are expedited per repair resource. For this non-linear non-convex integer programming problem, we develop a greedy heuristic and an algorithm based on decomposition and column generation. Both solution approaches perform very well with average optimality gaps of 2.38 and 0.27 percent, respectively, across a large test bed of industrial size. The possibility to expedite the repair of failed parts is effective in reducing stock investments with average reductions of 7.94 percent and even reductions up to 19.61 percent relative to the state of the art. We report on a case study at the Dutch Railways and show that our model leads to saving potentials of up to 50 percent compared to the current way of working.

Sara Ghazanfari
CWI
s.ghazanfari@cwi.nl
Superviso Rudesindo Núñez Queija

Title *The dynamic user equilibrium models under hypercongestion and travel uncertainties using Vickrey's bottleneck model*
Abstract Bottleneck traffic congestion has been the subject of many research studies in the transportation area. We first study an extension of the seminal Vickrey bottleneck model by assuming that there is uncertainty about the actual time of arrival to the bottleneck. Such uncertainties are the consequence of a wide variety of causes external to travellers' decisions, including actual departure times, weather circumstances and traffic delays. We show that with these travel time uncertainties a user equilibrium can neither be a pure Nash equilibrium, nor a mixed equilibrium with continuous density.
We then set out to describe a bottleneck model where the bottleneck capacity is negatively affected by the length of the queue and we show the existence of the user equilibrium in the presence of hypercongestion.

Alberto Giudici
Erasmus University Rotterdam
giudici@rsm.nl
Supervisors Tao Lu and Rob Zuidwijk
Co- author Clemens Tielen, Technical University of Munich

Title *Reliable cost-efficient flows in stochastic time-varying networks*
Abstract In the case of freight transport, network flows are an important and well-studied modelling tool. Real-world freight transport networks, however, are inherently stochastic (e.g., regarding travel times) and the available connections in the network usually vary over time. As the reliability of deliveries is becoming a main concern for transport operators, we are motivated to study the computation of reliable, cost-efficient flows in stochastic, time-varying networks. To the best of our knowledge, this problem has not been studied so far.
For the context of hinterland container transport, we define a model to address the problem of determining the value of adaptively routing flow in a capacitated network with stochastic, time-varying transit times. Our model features a combination of a minimum cost flow formulation, which minimizes capacity booking costs, and of a Markov Decision Process, which outputs a policy to route the flow on the booked capacity at maximum reliability. The goal is to minimize the total costs incurred for booking capacity on services prior to the execution of an adaptive routing policy that achieves a given minimum

level of reliability. In our study so far, we show how to rewrite the problem into a single MIP; furthermore, we devise a heuristic solution approach which performs a neighborhood search in the space of booked capacities by solving the MDP at each iteration. Our preliminary results show that the marginal amount of additional booked capacity depends both on the level of reliability and the amount of flow. Our next step is to study the structure of this problem.
Keywords hinterland container transport, reliability, adaptive routing, stochastic time-varying network

---

Lieke van der Heide
University of Groningen
l.m.van.der.heide@rug.nl
Supervisors Bram de Jonge and Dirk Pieter van Donk

Title *Minimizing traffic delays due to short-term infrastructure maintenance projects*
Abstract In the upcoming decades, renewing and maintaining infrastructures in the Netherlands will be an immense task that will temporarily reduce road capacity. Prior to maintenance projects, traffic is in a user-equilibrium where no driver can improve their travel time by changing their route. This equilibrium is disturbed when some roads become unavailable, causing traffic delays. These delays should be taken into account when a maintenance schedule is determined.
Most of the current literature on joint maintenance planning for underground infrastructures and road networks does not take traffic delays into account and only incorporates a reduced cost for simultaneous maintenance at the same location. When traffic delays as a consequence of maintenance activities are taken into account, only long-term projects are considered where traffic is assumed to be always in an equilibrium. This underestimates the traffic delays for short-term projects, since traffic slowly adjusts to the road disturbance before reaching a new equilibrium.
In this study, we consider a given set of short-term maintenance projects that disturb a road network and cause drivers to slowly adjust their routes according to a day-to-day traffic flow dynamic. Moreover, we consider two types of drivers that react to road network disruptions differently. Due to the complexity of the problem, we develop heuristic approaches for scheduling the maintenance projects, taking into account the resulting traffic delays. Infrastructure managers can use our approach to assess the benefit of jointly scheduling maintenance projects for multiple infrastructures and quantify the effect of resulting traffic delays.

Lucas van Krefeld
University of Amsterdam
l.r.vankrefeld@uva.nl
Supervisor Jan-Pieter Dorsman, Michel Mandjes and Onno Boxma

Title *Scaling analysis of an extended machine-repair model*
Abstract We consider an extension of the classic machine-repair model, where we explicitly model the fact that machines, apart from requiring service from a single repairer, also supply service themselves to products. Due to this dual role of the machines, the system exhibits an intricate relation between the processing rate of products and the performance of the repairer. To characterize this relation, we analyze this model under a Halfin-Whitt inspired scaling regime, where we amplify the arrival rate of products, the repair speed of the repairer and the number of machines appropriately. The resulting limiting stationary distribution is elegant, allows for a closed-form expression and provides intuition on the system's behavior resulting from the machines' dual role. With numerical results we illustrate the convergence, and assess under which conditions the limiting distributions lead to accurate approximations.

Roel Lambers
Eindhoven University of Technology
r.lambers@tue.nl
Supervisor Frits Spieksma

Title *Multi-league scheduling; how to schedule thousands of matches*
Abstract We consider the simultaneous scheduling of multiple sport leagues, involving thousands of matches, with interdependencies arising from teams in different leagues belonging to the same club. Teams from the same club share the same venue, which consists of a limited number of terrains (referred to as the club's capacity). The problem is to assign teams to Home-Away patterns while minimizing the total capacity violation over all clubs during the season. We show that this Multi-league Scheduling Problem (MSP) can be solved in polynomial time when each league consists of the same, even number of teams. Also, we introduce two generalizations of MSP: one where teams from the same club have to play according to the same Home-Away pattern, and one where club capacities differ throughout the season. We establish the borderline between easy and hard problems for each of these generalizations when varying the size of the league.

Weina Ma
Erasmus University Rotterdam
ma@ese.eur.nl
Supervisors Rommert Dekker and Chiel van Oosterom

Title *Admission control using real-time replenishment information in a make-to-stock system*
Abstract We consider a make-to-stock company's admission control and replenishment decisions with real-time status information on replenishment item in service control towers. The real-time status information is captured by a k-Erlang replenishment lead time. First we model the problem with patient customers as an infinite-horizon Markov decision process and minimize the total expected discounted cost. We prove that the optimal policy can be characterized using two thresholds: a base work storage level that determines when ordering takes place and an acceptance work storage level that determines when demand of customers is accepted. In the numerical study, we show that using real-time status information on replenishment item and adopting admission control can lead to significant cost savings. The cost savings are highest when the optimal admission threshold is a work storage level with a replenishment item halfway in process. This finding is different from the literature, where it is stated that the cost increase of ignoring real-time information is negligible under either the lost sales or the backordering case. Next we study the problem where customers are of limited patience. We find that the optimal admission policy is not always of threshold type. This is different from the literature which assumes an exponential production lead time.

Mirjam Meijer
Eindhoven University of Technology
m.s.meijer@tue.nl
Supervisors Ton de Kok and Willem van Jaarsveld

Title *Supply Chain Coordination with Service-Dependent Contract Termination*
Abstract In high-tech manufacturing, shortage of a single component from a supplier can lead to extremely costly delivery delays of a manufacturer's product. To avoid such shortages a manufacturer can incentivize her suppliers to develop sufficient costly capacity. To provide strong incentives, a manufacturer may develop a long-term relationship with her supplier. We model this relation as a repeated Stackelberg game in which the manufacturer is the natural leader. After each epoch, corresponding to one generation of the product, the manufacturer decides on her supplier for the next

product generation. We compare the case in which this decision takes into account performance of the current supplier to the case where it is not affected by performance, but purely relies on price considerations, and evaluate the effect on individual and supply chain profits. We show that, especially when the supplier has a high valuation of future profits, long-term relationships based on supplier performance can improve supply chain coordination.

Nemanja Milovanovic
Erasmus University Rotterdam
milovanovic@ese.eur.nl
Supervisor Rommert Dekker

Title *Stackelberg competition in liner shipping network design*
Abstract liner shipping company deploys a heterogeneous fleet of vessels in order to transport cargo from origin to destination in their service network, maximizing profit. Although several exact and heuristic approaches have been proposed for this problem, most of them assume that demand and market price are fixed regardless of how much capacity is offered for each shipping market. In addition, current literature ignores the effects of competition in liner service network design. The latter effect is particularly important, as in practice competition is fierce, with liner shipping companies ordering ever larger container ships in order to utilize economies of scale and offer lower shipping rates.
In order to capture the above, we propose a two-phase Stackelberg quantity competition model. In the first phase, the leader simultaneously constructs a service network and decides for each shipping market how many containers to transport. In the second phase, the follower takes similar actions, having observed the leader. This model is then solved as a bilevel program. Furthermore, we provide some insights and propose a method to solve this non-convex problem.

Maarten Otten
University of Twente
j.w.m.otten@utwente.nl
Supervisors Richard Boucherie and Aleida Braaksma

Title *Online routing for bed cleaning at nursing wards: a decomposed MDP approach*
Abstract Between the discharge of a patient and the admission of another, beds at nursing wards need to be cleaned. Beds can be cleaned at a central cleaning facility or at the wards. The latter is preferred to reduce the number of bed movements in the hospital. For this, a bed-cleaning team visits the wards at fixed times during the day. A high patient turnover and a lack of buffer capacity to store clean beds result in relatively small time windows in which beds should be cleaned. Due to the high variability of discharge times, a fixed cleaning schedule does not suffice. In this research we propose a Markov Decision Process (MDP) to model the routing of the bed-cleaning team in an online fashion. As a result of the curses of dimensionality this MDP is intractable. To approximate the MDP we decompose it into small-scale MDPs, in which a single ward optimizes its decisions given the decisions of the other wards. The interplay of the small-scale MDPs is modeled by a Non-Cooperative Stochastic Game (NCSG). We show that a best-response scheme results, after finitely many iterations, in a Nash-equilibrium of the NCSG. To reduce the gap between the optimal values of the MDP and the NCSG we extend the NCSG with several urgency levels. For this NCSG we show that a randomized best-response scheme converges with probability one to a Nash-equilibrium. We demonstrate the performance of our approach for a case study at a hospital.

Youri Raaijmakers
Eindhoven University of Technology
y.raaijmakers@tue.nl
Supervisors Sem Borst and Onno Boxma

Title: *Stability of Redundancy Systems with Processor Sharing*
Abstract We investigate the stability condition for redundancy-*d* systems where each of the servers follows a processor-sharing (PS) discipline. We allow for generally distributed job sizes, with possible dependence among the *d* replica sizes being governed by an arbitrary joint distribution. We establish that the stability condition is characterized by the expectation of the minimum of *d* replica sizes being less than the mean interarrival time per server. In the special case of identical replicas, the stability condition is insensitive to the job size distribution given its mean, and the stability threshold is inversely proportional to the number of replicas. In the special case of i.i.d. replicas, the stability threshold decreases (increases) in the number of replicas for job size distributions that are NBU (NWU). We also discuss extensions to scenarios with heterogeneous servers.

Victor Reijnders
University of Twente
v.m.j.j.reijnders@utwente.nl
Supervisors Johann Hurink, Marco Gerards and Gerard Smit

Title *Electricity Pricing Mechanism based on Losses Using Network Topology*
Abstract Pricing mechanisms for electricity network costs currently used for consumers do not provide incentives to decrease the stress on the electricity network. In this talk, we present a novel pricing mechanism for the network costs based on the losses caused by the transport of electricity in the low voltage (LV) network. This mechanism is based on the Shapley value. Firstly, we derive an explicit formulation of the Shapley value based on the given location of households in the LV network. Secondly, since the location of the households heavily influences the prices, we present a pricing mechanism that averages the Shapley value over all permutations of households on the different locations in the network. This results in a cost per household as if the household is on an average location in the network. This cost can be calculated efficiently if the households are connected to a single cable. Thirdly, an approximation of this value is presented which could be useful or even necessary for extensions to more complex electricity networks which consist of more than a single cable. Lastly, we show some considerations for the actual implementation of these pricing mechanisms.

Dennis Schol
Eindhoven University of Technology
c.schol@tue.nl
Supervisors Bert Zwart and Maria Vlasiou

Title *Large fork-join networks with nearly deterministic service times*
Abstract A typical property of high-tech manufacturing is that a large number of suppliers are involved, and are specialized in producing and delivering a very specific component of the final product. This study is motivated by modeling the delays in the emerging complex supply chain. Consider such a chain with many suppliers, each producing a specific component of the final product. In this system, the delay of the manufacturer is determined by the slowest supplier.
To model this delay, we propose a discrete time *N* server fork-join queueing network, in which each server represents a unique supplier. Each time step we have one arrival with a given probability *p*, and each server completes one service with probability *q*, independent of the other servers. These *p* and

q are close to 1, thus we have nearly deterministic arrivals and services. The aim of this study is to approximate the length of the largest of the *N* queues in the network.
We present a fluid limit and a steady-state result for the maximum queue length, as *N* goes to infinity. In order to get the fluid limit, we have to scale time with N^3 and space with *1/N*. For the steady-state result, we only have the spatial scaling of *1/N*. These two results have remarkable differences. The steady-state result depends on two model parameters, while the fluid limit only depends on one model parameter. Moreover, we see that the maximum queue length has a different scaling order in steady state compared to the fluid limit.
In order to prove these results, extreme value theory and diffusion approximations for the queue lengths are used. Since each queue has the same arrival process, queue lengths are dependent random variables, which makes it challenging to find convergence results of the maximum queue length. By giving upper and lower bounds on the maximum queue length, each having a small dependent and a large independent part, we are able to prove these convergence results. Finally, we do some numerical analysis to determine the phase transition between the fluid limit and the steady state result.

Céline Swennenhuis
Eindhoven University of Technology
c.m.f.swennenhuis@tue.nl
Supervisors Jesper Nederlof and Frits Spieksma

Title *Parameterized Complexity of Partial Scheduling*
Abstract We study a natural variant of scheduling that we call partial scheduling: In this variant an instance of a scheduling problem along with an integer k is given and one seeks an optimal schedule where not all, but only k jobs have to be processed.
We study the Fixed Parameter Tractability of partial scheduling problems parameterized by k for all variants of scheduling problems that minimize the makespan and involve unit/arbitrary processing times, identical/unrelated parallel machines, release/due dates, and precedence constraints. That is, we investigate if algorithms with runtimes of the type O*(f(k)) exist, where the O*(·) notation omits factors polynomial in the input size. We obtain a trichotomy by categorizing each variant to be either in P, NP-complete and Fixed Parameter Tractable by k, or W[1]-hard by k.
This presentation will focus on one of our main technical contributions, which is an O*(8k) time algorithm to solve instances of k-scheduling problems minimizing the makespan with unit job lengths, precedence constraints and release dates.

Eline Tsai
University of Twente
e.r.tsai@utwente.nl
Supervisors Richard Boucherie, Derya Demirtas, Yolanda de Rijke and Robert de Jonge

Title *Assignment of Clinical Chemistry Samples to Analyzers Lines to Reduce Sojourn Times*
Abstract In this study, we consider clinical chemistry laboratories using an automated pre-analysis unit for sample preparation and several automated analyzer lines for sample testing. In these systems, individual samples are already assigned to an analyzer at the start of pre-analysis. Prediction of the workload of the analyzers upon arrival of the samples in the analytical phase may improve analyzer selection with respect to sample sojourn times. To optimize analyzer selection, we have developed a queueing model that enables prediction of the sojourn time of samples in the pre-analytical and analytical phase.

Rogier Wuijts
Utrecht University
r.h.wuijts@uu.nl
Supervisors Marjan van den Akker and Machteld van den Broek

Title *An Improved Algorithm for Single-Unit Commitment with Ramping Limits*
Abstract The single-unit commitment problem (1UC) is the problem of finding a cost optimal schedule for a single generator given a time series of electricity prices subject to technical restrictions. We propose two efficient dynamic programming algorithms. For each timestep we keep track of a set of functions that represent the cost of optimal schedules until that timestep. We show that we can combine a subset of these functions by only considering their minimum. We can construct this minimum either implicitly or explicitly.
Experiments show both methods result in a significant speedup compared to state-of-the-art for piece-wise linear as well as quadratic generation cost. The time complexities of both methods seem to scale linear in the amount of timesteps. Therefore using these methods could lead to significant improvements for solving large scale unit commitment problems with Lagrangian relaxation or related methods that use 1UC as subproblem.

Mathijs van Zon
Erasmus University Rotterdam
vanzon@ese.eur.nl
Supervisors Remy Spliet and Wilco van den Heuvel

Title *Verifiable Stability in Collaborative Transport*
Abstract We consider a setting in which multiple companies collaborate in their transport operations. Each company has a set of delivery locations requiring a visit from a vehicle to satisfy demand. Although it is well-known that collaboration can lead to reduced transport costs, the remaining costs still need to be allocated to each company. A standard approach is to model this collaboration as a cooperative game, in which each company is referred to as a player and each set of players as a coalition. With each coalition a cost is associated, which represents the costs of transport operations. To allocate costs, it is standard to use a so-called core allocation. We provide a definition of the traditional core of a cooperative game and comment on the limitations of using a core allocation in practice. We introduce the concept of verifiable stability and provide an allocation method which is verifiably stable. We argue that our verifiably stable allocation method overcomes the limitations of a traditional core allocation, while in some sense upholding the concept of rationality on which core allocations are based. We will perform experiments to demonstrate the tractability of our approach and compare it to traditional core allocation methods.
