

Abstracts of the PhD presentations

Marleen Balvert

Tilburg University

m.balvert@uvt.nl

Supervisor Dick den Hertog

Title *Robust treatment plan optimization for prostate HDR brachytherapy incorporating delineation uncertainties*

Abstract High dose rate brachytherapy is a type of radiotherapy that is very suitable for the treatment of prostate cancer. A radioactive source is placed inside the tumor for a predetermined amount of time, the dwell time, and is then moved to the next position. The source locations and dwell times determine the radiation dose delivered to each part of the tumor as well as the surrounding organs at risk.

Mathematical optimization models are used to find the optimal dwell locations and times, where the aim is to deliver the prescribed dose to the tumor while limiting the dose to the organs at risk. The locations of the tumor and surrounding organs are used as an input for this model. For this purpose, a scan of the patient is made on which the medical physicist delineates the organs of interest. These delineations and thus the organ locations are subject to uncertainty. As a result, the treatment plan may be based on incorrect organ locations yielding the risk of insufficient dose to the target.

In this work, we develop a model that takes uncertainties in tumor delineations into account using robust optimization. Based on the delineations, scenarios of the true tumor location are generated. The model then optimizes for the worst case scenario.

Our method has been tested using the data obtained from three prostate cancer patients. Compared to current clinical practice, the results show a clear improvement in target coverage for most of the scenarios, without the need of increasing the dose to the organs at risk.

Thije van Barneveld

CWI

t.c.van.barneveld@cwi.nl

Supervisor Rob van der Mei, Sandjai Bhulai

Title *A Heuristic Method for Relevant Performance Measures in a dynamic Ambulance Management Model*

Abstract In life-threatening emergency situations, the ability of ambulance service providers to arrive at an emergency scene within a short period of time can make the difference between survival or death for the patient(s). In line with this, a service-level target related to the response time is commonly used. To realise such short response times, but still ensure running costs remain affordable, it is critical to efficiently plan ambulance services. This encompasses a variety of planning problems at the tactical, strategic and operational levels. A factor that further complicates this type of planning problems is uncertainty. The issue is that the planning methods currently available typically assume that ‘demand’ (in the context of ambulance services this would refer to call volumes and their geographical spread) and ‘supply’ (the availability of vehicles and ambulance personnel at the right time in the right place) parameters are known a priori. This makes these methods highly vulnerable to randomness or uncertainty, and the impacts this inevitably has on the broader planning process, namely inefficiencies, and higher costs. For ambulance services, the challenge is to develop new planning methods that are both scalable and robust against the inherent randomness associated with the service process, both in real and non-real time. In this talk, we study the Dynamic Ambulance Management (DAM) problem in which one tries to retain the ability to respond to possible future requests quickly when ambulances become busy. To this end, we need relocation actions for idle ambulances. Different performance measures related to response times can be incorporated in this model. We model the region of interest as an equidistant graph and we take into account the current status of both the system and the ambulances in a state. We do not require ambulances to return to a

base station, but they are allowed to idle at any node. This brings forth a high degree of complexity of the state space. Therefore, we present a heuristic approach to compute redeployment actions. We construct several scenarios that may occur one time-step later and combine these scenarios with each feasible action to obtain a classification of actions. We show that this heuristic policy performs well with respect to a policy often used in practice: the compliance table policy.

Joost Berkhout

VU University Amsterdam

j2.berkhout@student.vu.nl

Supervisor Bernd Heidergott

Title *Series Expansion Approximation for Multi-Chain Markov Chains*

Abstract In this talk we consider Markov chains which have multiple ergodic classes and transient states, so-called multi-chain Markov chains. We are interested in computing the ergodic projector of these multi-chains, where the ergodic projector of a Markov chain describes the asymptotic behavior of the chain in a concise way. If the ergodic classes and the transient states are known, efficient algorithms for computing the ergodic projector exist. In this talk we consider the case where the underlying structure of the multi-chain is unknown. We will explain that in computing the ergodic projector, the ergodic classes and transient states are identified.

We propose a series expansion approximation algorithm (SEA) to find an accurate approximation for the ergodic projector of a multi-chain. SEA has the advantage that the underlying structure of the multi-chain Markov chain does not have to be identified at the expense of an approximation. But as we will see in the convergence analysis this approximation error can be made arbitrarily small. Some numerical experiments will be presented to compare SEA with the straightforward power method. Possible applications are in Google's PageRank algorithm and in the identification of communities in social networks.

Gianmarco Bet

Eindhoven University of Technology

g.bet@tue.nl

Supervisor Johan van Leeuwen, Remco van der Hofstad

Title *Heavy-traffic analysis of transitory queues with diminishing populations*

Abstract We introduce a new queueing model that assumes a finite pool of customers. As time progresses, fewer customers can potentially join the system, causing the so-called *depletion of points* effect.

We are particularly interested in critically-loaded systems (close to 100% utilization), for which we derive stochastic-process limits using a martingales, functional limit theorems, and the Skorokhod mapping. One of the stochastic-process limits is a reflected Brownian motion with a parabolic drift, which also plays a key role in critical random graphs and epidemics.

Arjan Dijkstra

University of Groningen

a.s.dijkstra@rug.nl

Supervisor Kees Jan Roodbergen

Title *Optimal storage assignment in a manual-picking warehouse*

Abstract Order picking concerns the retrieval of products from storage locations in response to customer orders, and is one of the most time-critical processes in warehouses. In this paper, we focus on the combined effects of the routing methods and the storage location assignment on process performance. Routing methods determine the sequence in which locations are visited to retrieve any given customer order, while storage location assignment concerns the determination of locations

where products are to be stored. For four common routing methods, we give exact formulas for the average route length under any storage location assignment. These formulas subsequently serve as objective functions in an optimization problem to determine the best storage location assignment. Properties of optimal solutions are derived based on the route length formulas. Furthermore, we provide procedures for determining optimal or near-optimal solutions for storage location assignment, which rely on these optimality properties. Experiments underline the importance of optimization procedures for this problem by revealing patterns for storage assignment that have not been described in literature before.

Evelot Duijzer

Erasmus University Rotterdam

duijzer@ese.eur.nl

Supervisor Rommert Dekker and Willem van Jaarsveld

Title *Optimal Vaccine Allocation over multiple populations in the SIR model*

Abstract Vaccination is considered to be one of the most effective ways to prevent an epidemic. However, the size of vaccine stockpiles is often limited which complicates the allocation. Previous studies use simulation to compare different allocation policies, but these simulation results are not always easily understood. We analytically determine the optimal vaccine allocation over multiple non-interacting populations and derive structural results. We make use of a standard SIR epidemic model, which describes the evolution of an epidemic with differential equations.

The objective used is to minimise the total number of people infected and we analyse in what way the optimal allocation depends on the moment of vaccination. To obtain the results, we make extensive use of the analysis of implicit functions, optimisation and proofs of convexity. We present an optimal vaccination fraction: the vaccination fraction that gives the highest decrease in final size per dose of vaccine. Surprisingly, this vaccination fraction does not coincide with the critical vaccination coverage from literature. In order to reach the optimal vaccination fraction as close as possible, a subset of populations is determined in the optimal solution and the vaccine stockpile is allocated only to the populations in this subset.

Keywords: resource allocation, optimisation, vaccination, disease modelling, infectious diseases.

Martijn van Ee

VU University Amsterdam

m.van.ee@vu.nl

Supervisor Rene Sitters, Leen Stougie

Title *Approximation of a priori routing problems*

Abstract The field of *a priori* optimization is an interesting subfield of stochastic combinatorial optimization that is well suited for routing problems. In this setting, one has to construct a tour in the first stage and there is a probability distribution over active sets, vertices that have to be visited in the second stage. For a fixed first-stage tour, the second-stage tour on an active set is obtained by restricting the tour on the active set. In the *a priori* traveling salesman problem (TSP), the goal is to find a first-stage tour that minimizes the expected length of the second-stage tour. In the *a priori* traveling repairman problem (TRP), the goal is to find a tour that minimizes the expected sum of latencies. In this talk, I will give a brief overview of the results that were recently obtained for the *a priori* TSP and show that it matters how the probability distribution is represented. Moreover, I will discuss how we obtained the first constant approximation for the *a priori* TRP. To get this result, it suffices to give constant approximations for certain variations of the connected facility location problem. Finally, I will discuss some open problems and comment on some work in progress.

Jiwen Ge

Eindhoven University of Technology

j.ge@tue.nl

Supervisor Jan Fransoo, Dorothee Honhon

Title *Selling to Nanostores*

Abstract Nanostores, traditional small non-chain retailers prevalent in mega-cities of emerging markets, contribute a large proportion of revenue for many Consumer packaged Goods (CPG) manufacturers. These manufacturers send salespersons regularly to convince nanostores to order appropriately, by which we call sales effort. Nanostores have limited shelf space. Without sales effort, they place orders strongly based on previous periods' sales, by which we name as Sales Chasing. Our paper incorporates shelf space constraint and sales chasing behaviour. Distinguished in sales effort set up cost K being zero or positive, the paper builds two finite horizon Markov Decision Process models to trade-off the benefits and cost of sales effort. We categorize the analysis in difference cases and prove the corresponding optimal policies in general or under some mild technical conditions. Some relevant parametric results are also derived.

Qianru Ge

Eindhoven University of Technology

q.ge@tue.nl

Supervisor Geert-Jan van Houtum, Hao Peng

Title *Reliability Optimization for Critical Components with Uncertain Component Reliability in the Design*

Abstract We develop an optimization model to determine the reliability design of critical components in a system. Since the system is under a service contract, a penalty cost should be paid by the OEM when the total system down time exceeds a predetermined level, which complicates the evaluation of the life cycle costs. Furthermore, in the design phase for each critical component, all the possible designs are subject to uncertain component reliability. An efficient approximation method is proposed.

Sybren Huijink

Tilburg University

s.huijink@uvt.nl

Supervisor Goos Kant, Rene Peeters

Title *Which orders to outsource in the VRP with Private fleet and Common carrier*

Abstract In practice, many parcel transportation companies lower their costs by hiring carriers to deliver orders that cannot be served efficiently by their own fleet. The variant of the VRP which takes this outsource option into account is known as the Vehicle Routing Problem with Private fleet and Common carrier (VRPPC). The test instances in the literature for the VRPPC have relatively high outsource prices that serve as penalty costs for not having enough capacity. As a result, outsourcing is done out of necessity. To get more realistic instances we created new pricing schemes with lower outsourcing prices such that outsourcing is done out of efficiency, i.e., it might be more profitable to leave some capacity unused. To be able to get efficient results, we created an Adaptable Variable Neighborhood Search (AVNS) which has several totally different moves to tackle the problems on all instances. It turns out that our AVNS is highly competitive compared with the other heuristics in the literature for the VRPPC with original prices.

In practice, the amount of transport optimization solutions supporting the outsource option is also quite limited. Therefore, we investigated how one could estimate in a simple way which orders to outsource and determined what the loss would be if one then solves the remaining VRP problem. For this we compare existing methods and propose a new method, delivering better results.

Bram de Jonge

University of Groningen

b.de.jonge@rug.nl

Supervisor Ruud Tuenter, Tiedo Tinga

Title *On the Benefits of Condition-Based Maintenance over Time-Based Maintenance*

Abstract Due to ongoing automation of production processes and increasing reliance on expensive production equipment, the importance of effectively planned and performed preventive maintenance activities is growing. Many firms still apply 'traditional' time-based maintenance (TBM) strategies that are easy to implement as only the time that a unit is in service has to be recorded. However, a substantial remaining useful life is wasted if the machine is still in reasonable condition when preventive maintenance is performed, and a breakdown might occur if it happens to deteriorate faster than expected. Due to the increasing technical possibilities to monitor, store, and analyze conditions, condition-based maintenance (CBM) is gaining popularity.

CBM generally results in more effectively scheduled preventive maintenance, and, in the ideal case, in preventive maintenance that is performed just before failure. However, the relative performance of CBM is strongly dependent on the characteristics of the situation. In this study, we provide general insights on how the cost benefit of CBM compared with TBM is influenced by the behavior of the deterioration process, the severity of failures, the setup time required to perform preventive maintenance, the accuracy of the measured condition information, and the uncertainty in the deterioration level at which failure occurs.

The obtained insights are important in practice, because CBM should only be applied if its benefits over TBM outweigh the efforts and costs that are required to implement CBM. The requisites to switch from time-based to condition-based maintenance include condition monitoring equipment and software to store, analyze, and initiate maintenance actions.

Bart Kamphorst

CWI

b.kamphorst@cwi.nl

Supervisor Nikhil Bansal, Bert Zwart

Title *Achievable Performance of Blind Scheduling Policies in Heavy Traffic*

Abstract It is well-known that SRPT is an optimal scheduling policy for minimizing the average sojourn time in G/G/1 queues. Unfortunately, in practice one often lacks information on individual job sizes and hence a blind policy is required. Scheduling literature has shown strong bounds on the performance of blind scheduling policies in deterministic environments when compared to SRPT. Similar results for stochastic environments are scarce. In this presentation I will show how the strong bounds on blind policies in deterministic environments can be exploited in order to obtain bounds in heavy traffic stochastic environments.

In particular, I will introduce the Randomized Multi-Level Feedback (RMLF) algorithm and compare it to SRPT. RMLF is a scheduling policy that tries to mimic the behavior of SRPT without knowing job sizes. After a brief introduction of both policies I will address the question how the two scheduling policies compare in heavy traffic. Insights and techniques from scheduling theory, game theory and queueing theory are combined in order to answer this question. The stochastic analysis builds on a characterization of the cycle maximum for heavy-tailed job sizes in heavy traffic. One of our work's contributions is unification of a well-known result in this area, which relates the cycle maximum tail distribution to the job size tail distribution.

Thijs van der Klauw

University of Twente

t.vanderklauw@utwente.nl

Supervisor Johann Hurink, Gerard Smit

Title *Local device scheduling for demand side management using two types of steering signals*

Abstract Within power grids, supply and demand need to be matched at all times. This is traditionally done by using flexibility in supply to match demand. However, this methodology is less viable in future power grids due to an increasing share of inflexible generation from renewable sources such as wind and sun. Thus flexibility to match supply and demand needs to be sought elsewhere.

It is expected that in the future an increasing amount of flexibility is available on the demand side of electricity, mainly in the form of appliances that store energy in some form or manner. Examples of such appliances are electrical vehicles, heat pumps combined with heat buffers and fridges. Demand Side Management (DSM) methodologies often attempt to schedule these appliances locally based on a steering signal sent by a global controller. The local controllers use these steering signals to determine an optimal schedule for their appliances.

The local schedule generated by the local controller is thus an important step within the realization and analysis of DSM methodologies. To this end we study a two-fold steering signal consisting of both a time-varying price and a target profile. We model the local minimization objective as a weighted sum of the total cost of the consumed energy and the squared deviation from the target profile. The minimization is done subject to local constraints implied by the appliance. We study the structure of the derived optimization problems and use them to obtain efficient algorithms that solve the optimization problems to optimality.

Vincent Kreuzen

Maastricht University

v.kreuzen@maastrichtuniversity.nl

Supervisor Tobias Harks

Title *Basic Properties for High Multiplicity Scheduling with Switching Costs*

Abstract We study several variants of the single machine capacitated lot sizing problem with sequence-dependent setup costs and product-dependent inventory costs. Here we are given one machine and k types of products that need to be scheduled, each associated with a constant demand rate, production rate $p(i)$ and inventory costs per unit. When the machine switches from producing product i to product j , setup costs $c(i,j)$ are incurred. The goal is to find a schedule such that demand is met at all times and the average per-time-unit costs are minimized. This can be seen as lifting a conventional scheduling problem to its more general high-multiplicity counterpart where there are only a few job types, but each with a high multiplicity. This severely increases the complexity of the problem.

We distinguish three cases. In the continuous case the machine can switch products at any time and it can produce at most $p(i)$ units of product i . In the discrete case the machine can only switch products at the end of every unit of time (e.g. a day) and it can produce at most $p(i)$ units of product i . In the fixed case the machine can only switch products at the end of every unit of time and if it produces product i during that unit of time, it has to produce exactly $p(i)$ units.

We characterize feasible instances and prove additional structural properties. We prove the decision variants of the cases are in P and we provide an algorithm which outputs a polynomial-sized representation of an optimal schedule. We characterize optimal solutions for few products, where the fixed case for $k=1$ is already non-trivial.

Baoxiang Li

Eindhoven University of Technology

b.li@tue.nl

Supervisor Dmitry Krushinsky, Tom van Woensel, Hajo Reijers

Title *Rebalancing in Bike/Car Sharing System: Mathematical Formulations and a Mathheuristic Approach*

Abstract This paper deals with a pickup and delivery problem motivated by bicycle sharing systems with demand ranges. In a bicycle sharing system, the stations are required to have an inventory of bicycles within given lower and upper bounds, based on historical user data. Some stations may have

higher demand than others. If no action is taken, the inventory in these stations may rapidly reach a bound, thus preventing other passengers from picking up or dropping off bikes. A solution is to use vehicles to transport bikes from full stations to stations with shortages to balance the network. We formally define the problem and present a mathematical formulation. This formulation, however, is complex to get an exact solution. We propose a heuristic method to solve this problem, which includes three steps: (i) tabu search to get an upper bound, (ii) Lagrangean to get a lower bound, (iii) improve the final solution based on tabu search or branch and cut algorithm. The method provides a good lower bound in a reasonable time. We thus believe that our approach is suitable for practical implementation in bike sharing systems. The proposed heuristic can be applied to other vehicle routing problems, as well as to other sharing systems.

Marieke Musegaas

Tilburg University

m.musegaas@uvt.nl

Supervisor Peter Borm, M. Quant

Title *Step out - Step in Sequencing Games*

Abstract In a one-machine sequencing situation there is a queue of players in front of a single machine, each with one job to be processed. A cooperative game that arises from this situation is called a sequencing game. The common assumption in many sequencing games is that two players of a certain coalition can only swap their positions if all players between them are also members of the coalition. This assumption is too restrictive because there may be more reorderings possible which do not hurt the interests of the players outside the coalition. Relaxed sequencing games arise by relaxing this assumption. We introduce a new class of relaxed sequencing games, the class of Step out - Step in (SoSi) sequencing games. This relaxation is intuitive from a practical point of view, because any player within a coalition is allowed to step out from his position in the processing order and to step in at any position later in the processing order. We provide a polynomial time algorithm to determine the coalitional values of a SoSi sequencing game, i.e., an algorithm determining for every possible coalition an optimal admissible processing order. This algorithm works in a greedy way in the sense that every player is moved to the position giving the highest cost savings at that moment. Moreover, every player is considered in the algorithm exactly once and every player is moved at most once. Additionally, we prove that every SoSi sequencing game has a non-empty core.

Martijn Onderwater

CWI/VU University Amsterdam

m.onderwater@cwi.nl

Supervisor Rob van der Mei

Title *Throughput analysis of the 802.15.4 MAC for sensor networks*

Abstract For this research we look at a sensor network containing a number of sensor nodes competing for access to the wireless channel. This competition is described by the 802.15.4 MAC protocol, and causes alternating busy periods (a packet transmission) and silent periods (no packet transmission) on the channel. We are interested in the throughput of the wireless channel, i.e., the fraction of time that the channel is busy with a packet transmission. It is our intention to analyse the throughput mathematically, and to verify it using both simulations in OPNET and experiments with a physical sensor network. In this presentation we explain our aim in more detail, and illustrate the progress so far.

Tim Oosterwijk

Maastricht University

t.oosterwijk@maastrichtuniversity.nl

Supervisor Tjark Vredevelde

Title *A logarithmic approximation for matroid congestion games*

Abstract We study the problem of computing a social optimum (minimum cost solution) in matroid congestion games, where the strategy space of every player consists of the set of bases of a player-specific matroid defined on the ground set of resources. For general non-decreasing cost functions we devise an Hrk-approximation algorithm, where rk is the sum of the ranks of the player-specific matroids and Hrk denotes the rk th harmonic number. The main idea of our algorithm is to iteratively increase resource utilization in a greedy fashion and to invoke a polynomial covering oracle that checks feasibility of every computed resource utilization. The approximation guarantee is best possible up to a constant factor given that for singleton congestion games, there exists a logarithmic hardness result in the number of players. Our result settles an open problem of Ackermann et al. (H. Ackermann, H. Röglin and B. Vöcking, On the impact of combinatorial structure on congestion games, J. ACM 55(6): 25:1-25:22, 2008, Section 2.2), where the complexity of computing a socially optimal solution for matroid congestion games with non-decreasing cost functions is considered.

Krzysztof Postek

Tilburg University

k.postek@uvt.nl

Supervisor Dick den Hertog, Bertrand Melenberg

Title *Multi-stage adjustable robust mixed-integer optimization via iterative splitting of the uncertainty set*

Abstract Multiperiod problems where some of the parameters are uncertain and later-period decisions are integer are common, e.g., in inventory management, energy production, or manpower management. Key difficulty of applying Adjustable Robust Optimization in such problems, i.e., finding the best decisions under the assumption that the parameters take always their worst possible values, lies in providing adjustability of later decisions (giving better objective values) while keeping the solution time of the problem down. Both of these tasks grow even more difficult in the mixed-integer case. We propose a methodology for constructing decision rules for integer and continuous decision variables in multiperiod robust linear optimization problems. We show that by iteratively splitting the parameters' uncertainty set into smaller subsets one can differentiate the later-period decisions based on the revealed uncertain parameters. At the same time, the complexity of the robust problem stays at the same level as for the static robust problem. We provide theoretical results how to split the uncertainty set by identifying sets of uncertain parameter scenarios to be divided for an improvement in the worst-case objective value. Based on this theory, we propose several splitting heuristics. Numerical examples entailing a capital budgeting and a lot sizing problem illustrate the advantages of the proposed approach.

Frans de Ruiter

Tilburg University

fjctderuiter@gmail.com

Supervisor Dick den Hertog

Title *Robust optimization of uncertain multistage inventory systems with inexact data in decision rules*

Abstract In production-inventory problems customer demand is often subject to uncertainty. Therefore, it is challenging to design production plans that satisfy both demand and a set of constraints on e.g. production capacity and required inventory levels. Adjustable robust optimization (ARO) is a technique to solve these dynamic (multistage) production-inventory problems. In ARO, the decision in each stage is a function of the data on the realizations of the uncertain demand gathered from the previous periods. These data, however, are often inaccurate; there is much evidence in the information management literature that data quality in inventory systems is often poor. Reliance on data "as is" may then lead to poor performance of ARO, or in fact of any "data-driven" method. In this paper, we remedy this weakness of ARO by introducing a methodology that treats past data itself as an

uncertain model parameter. We show that computational tractability of the robust counterparts associated with this extension of ARO is still maintained. The benefits of our new approach are demonstrated by a numerical test case.

Jaron Sanders

Eindhoven University of Technology

jaron.sanders@tue.nl

Supervisor Johan van Leeuwen, Sem Borst

Title *Optimal Admission Control for Many-Server Systems with QED-Driven Revenues*

Abstract We consider Markovian many-server systems with admission control operating in a QED regime, where the relative utilization approaches unity while the number of servers grows large, providing natural Economies-of-Scale. In order to determine the optimal admission control policy, we adopt a revenue maximization framework, and suppose that the revenue rate attains a maximum when no customers are waiting and no servers are idling. When the revenue function scales properly with the system size, we show that a nondegenerate optimization problem arises in the limit. Detailed analysis demonstrates that the revenue is maximized by nontrivial policies that bar customers from entering when the queue length exceeds a certain threshold of the order of the typical square-root level variation in the system occupancy. We identify a fundamental equation characterizing the optimal threshold, which we extensively leverage to provide broadly applicable upper/lower bounds for the optimal threshold, establish its monotonicity, and examine its asymptotic behavior, all for general revenue structures. For linear and exponential revenue structures, we present explicit expressions for the optimal threshold.

Loe Schlicher

Eindhoven University of Technology

l.p.j.schlicher@tue.nl

Supervisor Marco Slikker, Geert-Jan van Houtum

Title *Pooling of services with unavailability*

Abstract In this paper, we consider an environment wherein several independent service providers, each subject to unavailability, can collaborate by pooling their services. We examine the allocation of the collective cost savings for such pooled situation by studying an associated cooperative game. First, we prove non-emptiness of the core, total balancedness and, under some conditions, convexity. Then, four allocation rules will be introduced and we will investigate whether they satisfy interesting properties such as monotonicity and symmetry. Finally, we will also investigate whether the payoff vectors resulting from these allocation rules are core members.

Harwin de Vries

Erasmus University Rotterdam

hdevries@ese.eur.nl

Supervisor Albert Wagelmans, Joris van de Klundert

Title *The roadside Healthcare Facility Location Problem*

Abstract Providing African truck drivers with adequate access to healthcare is an effective way to reduce the burden and the spread of HIV and other infectious diseases. Therefore, NGO North Star Alliance builds a network of healthcare facilities along major African trucking routes. Choosing the locations of new facilities presents novel and complex optimization problems. In this talk, we consider a general design problem: the Roadside Health Care Facility location Problem (RHFLP). RHFLP entails to select locations for new facilities and to choose for each of these facilities whether or not to add healthcare services for HIV, STIs, Tuberculosis, and/or Malaria to the standard health service

package. The objective combines the maximization of the truck driver patient volume at these facilities and the maximization of the extent to which the truck drivers have continuous access to the needed health service packages. We present three measures for continuous access to health services by mobile patients and integrate these measures in a mixed-integer programming formulation for RHFLP. Moreover, we prove the RHFLP to be strongly NP-hard and derive analytical results for the worst-case effects of impreciseness in the input data. We show how large scale real life problem instances can be solved, presenting numerical experiments for the North-South corridor network (Southern and Eastern Africa), and discuss policy implications.

Jianzhe Zhen

Tilburg University

j.zhen@uvt.nl

Supervisor Dick den Hertog

Title *Computing the Maximum Volume Inscribed Ellipsoid of a Polytopic Projection*

Abstract This paper introduces a method for computing the maximum volume inscribed ellipsoid and k -ball of a projected polytope. It is known that deriving an explicit description of a projected polytope is NP-hard. By using adjustable robust optimization techniques, we construct a computationally tractable method that does not require an explicit description of the projection. The obtained centers of the projected polytope are considered as the robust solutions, e.g., for design centering problems. We perform numerical experiments on a simple polytope and a color tube design problem. The color tube design problem demonstrates that our solutions are much more robust than the nominal approach with a slight compromise on the objective value. Some other potential applications include ellipsoidal approximations to polytopical sets, nominal scenario recovery, and cutting-plane method.
