

Abstracts of the PhD presentations

Abhishek

University of Amsterdam

abhishek@uva.nl

Supervisors Onno Boxma, Michel Mandjes, Sindo Núñez Queija, Marko Boon

Title *Waiting time and heavy-traffic analysis of $M^X/G/1$ type queuing models with dependent service durations*

Abstract We consider a single-server queue with N types of services, in which, customers arrive according to a batch Poisson process. The service times of customers have general distribution functions and are correlated. The correlations are due to the different service types, which form a Markov chain that itself depends on the sequence of service lengths. In addition, the first customer in a busy period has a different service time distribution than regular customers served in the busy period. Our work is motivated by an application of the model to road traffic, where a stream of vehicles on a minor road merges with a stream on a main road at an unsignalized intersection such that the merging times of two subsequent vehicles on the minor road are dependent.

Based on the results from a previous paper on the steady-state distribution of the queue length, we derive the waiting time and sojourn time distributions. The waiting times and sojourn times of customers depend on their positions in the batches, as well as on the type of service of the first customers in their batches. We also determine the heavy-traffic distribution of the scaled stationary queue length.

Angelos Aveklouris

Eindhoven University of Technology

a.aveklouris@tue.nl

Supervisors Maria Vlasiou, Bert Zwart

Title *A Stochastic Resource-Sharing Network for Electric Vehicle Charging*

Abstract We consider a distribution grid used to charge electric vehicles subject to voltage stability and various other constraints.

We model this as a class of resource-sharing networks known as bandwidth-sharing networks in the communication network literature. Such networks have proved themselves to be an effective flow-level model of data traffic in wired and wireless networks. We focus on resource sharing networks that are driven by a class of greedy control rules that can be implemented in a decentralized fashion. For a large number of such control rules, we can characterize the performance of the system, subject to voltage stability constraints, by a fluid approximation. This leads to a set of dynamic equations that take into account the stochastic behavior of cars. We show that the invariant point of these equations is unique and can be computed by solving a specific ACOPF problem, which admits an exact convex relaxation. For the class of weighted proportional fairness control, we show additional appealing properties under the linearized Distflow model, such as fairness, and a product form property of the stochastic model.

Annelieke Baller

VU Amsterdam

a.c.baller@vu.nl

Supervisors Wout Dullaert, Leen Stougie

Title *Integrating transportation and inventory management in cash supply chains*

Abstract The Dutch cash supply chain is characterized by a vendor-managed inventory setting in which a supplier (Geldservice Nederland) determines the timing and size of replenishments for its customers. The actual transport is performed by so-called cash-in-transit companies operating armored cars. To efficiently organize the cash supply chain, inventory management and transport decisions have to be optimized which poses several research opportunities.

First, we extend the Dynamic-Demand Joint Replenishment Problem (DJRP). In the DJRP one assumes that the supplier pays a fixed fee for replenishing a customer which often occurs if the supplier outsources transportation. Hence, there is no incentive for the supplier to schedule replenishments for nearby customers in the same period. To lower costs for both parties, we extend the traditional DJRP to the DJRP with Approximated Transportation Costs (DJRP-AT) by taking transportation considerations into account. A solution approach for the DJRP-AT based on Branch-and-Cut-and-Price is validated using test instances from the literature and on a real-life case.

Second, we study an extension of the Inventory Routing Problem (IRP). In the IRP, inventory replenishments and the routes of the delivery vehicles are optimized simultaneously. In large cities ATMs are often located in close proximity. This provides the opportunity to redirect a consumer that wants to withdraw money from one ATM to another in case of a cash shortage, hence, to move demand between ATMs. To best of our knowledge, the possibility of redirecting consumers to lower total inventory routing costs is new to the operations research literature and industry practice. We formulate the IRP with Demand Moves in which demand of an ATM can (partially) be satisfied by a nearby ATM. We propose a Branch-and-Price-and-Cut solution approach which is evaluated on problem instances derived from the literature and industry practice.

Mihail Bazhba

CWI

M.Bazhba@cw.nl

Supervisor Bert Zwart

Title *Sample-path large deviations for Lévy processes and random walks with Weibull increments*

Abstract While sample path large deviations with light tailed increments can be obtained due to the existence of the moment generating function, the number of similar results for heavy tails are scarce. We consider sample path large deviations for a Lévy process with heavy tailed Weibull increments. We develop a proof approach based on an appropriate representation, a suitable normalization, Bryc's inverse lemma, and the use of concentration inequalities. Our main results include an extended form of an LDP (large deviations principle) in the J1 topology, and a full LDP in the M1' topology, improving on a result in the L1 topology due to Gantert (1998). The rate function can be represented as the solution of a quasi-variational problem. The sharpness and applicability of these results are illustrated by a counterexample proving nonexistence of a full LDP in the J1 topology, and an application to the buildup of a large queue length in a queue with multiple servers.

Thomas Bosman

VU University of Amsterdam

Thomas.bosman@vu.nl

Supervisors Neil Olver, Leen Stougie

Title *Complexity of the Masked VPN Problem - Can We Generalize The VPN Theorem?*

Abstract The VPN problems cover an interesting class of optimization problems with the following interpretation. Given is a set of terminals with uncertain or time varying demand for sending data

traffic to each other through some network. We are required to pre-specify routing paths to be used between terminals. Our goal is then to minimize the cost of the capacity needed to make any possible demand pattern feasible to route. In the masked VPN problem, the set of feasible demand patterns is given by the set of fractional matchings in an auxiliary 'mask' graph, defined on the terminals.

The masked VPN problem generalizes a number of well-known VPN problems, like the symmetric and asymmetric variants as well as fundamental problems like Steiner tree. Hence, the problem is APX-hard, while it's an interesting open problem to beat a logarithmic approximation factor (which is achieved by using tree embeddings).

In this talk we look at potential generalizations of the celebrated VPN Theorem, which says that an optimal solution to the symmetric VPN problem has a sufficiently simple topology, that we can essentially 'guess' the optimal solution in polynomial time. We show similar structure theorems hold for several classes of masked VPN problems. Apart from drawing non-trivial links with existing literature (in particular, the 0-extension problem), these results form the basis for both exact and $O(1)$ -approximation (polynomial time) algorithms.

Joint work with Neil Olver. The results discussed in this talk appeared in "Exploring the tractability of the capped hose model", ESA 2017.

Mark van der Boor

Eindhoven University of Technology

m.v.d.boor@tue.nl

Supervisors Sem Borst, Johan van Leeuwen

Title *Scalable Load Balancing*

Abstract We consider a queueing system where jobs arrive at a central dispatcher. Every job needs to be sent to one of many identical servers where it will be served for an exponential time. The concept of (intelligently) dispatching jobs to servers is known as load balancing and we consider the system as the number of servers (as well as the arrival rate of jobs) tends to infinity. Many schemes exist; random assignment, Join-the-Shortest-Queue, power-of-d etc. While Join-the-Shortest-Queue achieves minimal wait, many communication between the dispatcher and the servers is needed, to be able to know which server has the shortest queue.

In this talk we will focus on strategies, for which the communication between the dispatcher and the servers is limited. Specifically, a constant number of queue lengths can be communicated per arriving job (on average). We introduce memory to the dispatcher, after which we introduce schemes that make clever use of this additional memory. We will analyze these schemes with fluid limits and show that they outperform many well-known load balancing schemes, in terms of both mean waiting time and communication overhead

Michiel uit het Broek

University of Groningen

a.j.uit.het.broek@rug.nl

Supervisors Ruud Tuentjer, Bram de Jonge, Jasper Veldman

Title *Condition-based production planning for systems with production-dependent deterioration*

Abstract Many large plants like power plants and refineries, commonly use so-called turnaround maintenance policies. In such policies, the entire system is shut down for a certain period and the whole system is maintained at once. This allows for the maintenance activities to be clustered and planned long in advance, thereby minimizing system downtime as well as logistics costs. However, the time between consecutive turnarounds is often large and machines may deteriorate faster than expected. In such situations, an interesting question is whether it can be profitable to reduce

production rates in order to avoid the need for maintenance before a turnaround.

The current maintenance literature typically assumes that machines always produce at the maximum production rate and that we therefore cannot influence the deterioration rate. However, there are many real life situations where we can adjust production rates. For example, wind turbines can decelerate which results in lower production rates as well as reduced deterioration rates. In this study, we consider a single unit with the ability to adjust production rates based on condition information. We conclude that the flexibility to operate with different production rates reduces the probability of a failure while the expected total production increases.

Thomas Breugem

Erasmus University Rotterdam

breugem.thomas@gmail.com

Supervisor Dennis Huisman

Title *Analyzing the Trade-Off between Fairness and Attractiveness in Crew Rostering*

Abstract In this paper, we analyze the trade-off between perceived fairness and perceived attractiveness in crew rostering. First, we introduce the Fairness-oriented Crew Rostering Problem. In this problem, attractive cyclic rosters have to be constructed, while respecting a pre-specified fairness level. Then, we propose a flexible mathematical formulation, and develop an exact Branch-Price-and-Cut solution method. Finally, we demonstrate the benefit of our approach on practical instances from Netherlands Railways, the largest passenger railway operator in the Netherlands. We are able to construct the explicit trade-off curve between fairness and attractiveness, and show that a sequential approach can lead to suboptimal results. In particular, we show that focusing solely on fairness leads to rosters that are disproportionately less attractive. Furthermore, this decrease in attractiveness is heavily skewed towards the most flexible employees. Thus, in order to generate truly fair rosters, the explicit trade-off between fairness and attractiveness should be considered.

Teun Janssen

Delft University of Technology

t.m.l.janssen@tud.nl

Supervisor Leo van Iersel

Title *Scheduling in the Photolithography Bay*

Abstract The production process in a semiconductor factory is a complex one. The wafer, which contains the chips, will visit different production bays multiple times during its production cycle. The expensive photo-lithography equipment often are the bottleneck of the production line. Hence, the overall performance of the wafer fab can be improved by raising the equipment throughput on these tools. Photolithography is a process to transfer the geometric pattern of a chip-design onto a wafer. This is done by putting light through a reticle (or mask) onto the production wafer. This reticle contains the geometrical pattern of the computer chip.

In this talk, we will look at scheduling problems for the photolithography bay. In European factories, this bay contains many different machines and the products are very diverse. Furthermore, there is only one reticle of every kind in the whole factory, thus products that share a reticle cannot be produced at the same time. This translates to a mathematical problem, which is a special case of scheduling with resource constraints. We will look at this problem and its properties and pose the major open questions

Madelon de Kemp

University of Amsterdam

madelon-de-kemp@hotmail.com

Supervisors Michel Mandjes, Neil Olver

Title *Performance of the smallest-variance-first rule in appointment sequencing*

Abstract In appointment scheduling problems, the deterministic arrival times of patients in a single-server queue should be determined. This should be done such that there is both little idle time and little waiting time. Part of this problem is the sequencing problem, in which the order in which different patients arrive should be determined.

The so-called smallest-variance-first rule, which sequences patients in order of increasing variance of their service durations, is often found to perform better than other heuristics. Under some simplifying assumptions, we will consider the performance of the smallest-variance-first rule. By comparing an upper bound on the expected waiting times under this rule with a lower bound valid for any sequence, we will be able to bound the relative difference in performance between the smallest-variance-first rule and the optimal sequence.

We also consider the limit as the number of patients tends to infinity. We show that the relative difference in performance approaches 1, thus proving that the smallest-variance-first rule is asymptotically optimal.

Joost de Kruijff

Eindhoven University of Technology

j.t.d.kruijff@tue.nl

Supervisor Cor Hurkens

Title *Integer variables in low-volume production planning models*

Abstract Production planning aims to coordinate the release of materials and capacity such that known or forecasted demand is met as good as possible. Both in literature and in practice usually linear programming models are applied in a rolling horizon approach. The majority of the literature considers high-volume settings and uses continuous production quantity variables in the models. For low-volume environments rounding these continuous production quantities results in poor plans. Therefore, in low volume production planning integer production quantities are required.

We discuss three production planning models in a low-volume setting: a simple production planning model, a lot sizing model and a model specific for high-tech low-volume industries with varying capacity availability. In these models integer production quantity variables imply the inventory variables to be integer. In two of these models also the opposite is true. This enables us to choose which variables are integer in the model: the inventory variable, the production quantity variable or both. We show that this choice has a big impact on the solvability of all three models.

Bart Litjens

University of Amsterdam

bart_litjens@hotmail.com

Supervisor Lex Schrijver

Title *Bounds in coding theory using semidefinite programming*

Abstract In this talk I will present some results on new upper bounds for the maximum cardinality $A_q(n,d)$ of a code of length n over an alphabet with q letters and with minimum distance at least d . Using semidefinite programming, an upper bound is obtained which is stronger than the Delsarte linear programming upper bound. By symmetry of the problem, representation theory can be applied to reduce to a semidefinite programming problem with order bounded by a polynomial in n . The same ideas can be applied to upper bound the maximum cardinality of a mixed binary/ternary code (i.e., with a fixed number of binary and ternary coordinates), and a constant weight code (i.e.,

with a fixed number of nonzero coordinates) with minimum distance at least d . I shortly discuss the first application. This talk is based on joint work with Sven Polak and Alexander Schrijver.

Mayank Saxena

Eindhoven University of Technology

m.mayank@tue.nl

Supervisors Onno Boxma, Rudesindo Núñez Queija, Stella Kapodistria

Title *A cooperative random access network under JSQ policy*

(Joint work with Stella Kapodistria and Ioannis Dimitriou)

Abstract In this paper, we consider a discrete time relay-assisted random access wireless network. This network is composed of a source user that transmits packets to a common destination node, with the cooperation of two relay nodes. We assume an Early Arrival System, and i.i.d. Bernoulli packet arrivals at the beginning of each time slot under the 'join the shortest queue' policy. At the end of each slot, both relays attempt to transmit a packet with given probabilities, independently of each other. If both relays attempt transmission at the same slot, there is a collision, and both packets have to be re-transmitted in a later slot.

For this model, we compare three different methods for computing the steady-state joint queue length distribution. First we show that this distribution can be written as an infinite sum of product form solutions, which satisfy a recurrence relation. Furthermore, we compute the steady-state joint queue length distribution using the power series method, as well as the truncation of state space. Finally, we present the numerical comparison of these three approaches to give a better understating of which method performs better than the others in terms of accuracy and computation time.

Anna Oblakova

University of Twente

a.oblakova@utwente.nl

Supervisors A. Al Hanbali, Richard Boucherie, Jan Kees van Ommeren, Henk Zijm

Title *Queueing in traffic networks*

Abstract In this talk, we present several discrete-time models, which describe queues at the traffic-light intersections. We begin with an isolated intersection with fixed-cycle control, i.e., each lane has fixed green and red times. This system was analysed in 1964 by Darroch [1], who showed that the probability generating function of the queue length can be represented as a rational function with several unknown probabilities. These unknown probabilities can be found using the roots of a characteristic equation. However, the root-finding procedure can be quite time-consuming and error-prone. For this model, we show that an elegant closed-form solution exists, which is root-free.

Analysis of traffic networks poses new challenges, as the output of one intersection becomes the input to a downstream intersection. To tackle the multidimensionality of the problem, we assume independence between cycles, and focus on one lane at a time. We also consider a more realistic departure process. For this model, we derive the probability generating function of the queue length, which, as before, can be represented as a rational function. In this case, however, we were not able to find a root-free solution. Despite that, we show that our model quite accurately predicts the results obtained using traffic microsimulation suite SUMO [2], while being more than 200 times faster.

[1] J. Darroch, On the traffic-light queue, *The Annals of Mathematical Statistics* 35 (1) (1964) 380-388

[2] D. Krajzewicz, J. Erdmann, M. Behrisch, and L. Bieker, Recent development and applications of SUMO - Simulation of Urban MObility. *International Journal On Advances in Systems and*

Sven Polak

University of Amsterdam
sven_polak@hotmail.com
Supervisor Lex Schrijver

Title *Uniqueness of codes using semidefinite programming*

(based on joint work with Andries Brouwer, <https://arxiv.org/abs/1709.02195>)

Abstract Fix a natural number n . A word is an element v in $\{0,1\}^n$. For two words u, v in $\{0,1\}^n$, their (Hamming) distance $d_H(u,v)$ is the number of positions in which they differ. The weight of a word v in $\{0,1\}^n$ is the number of ones in v . A binary constant-weight code C is a subset of $\{0,1\}^n$ in which all elements have a fixed weight w . The minimum distance of a code C is the minimum of $d_H(u,v)$ over all distinct codewords u, v in C .

For natural numbers n, d, w , let $A(n, d, w)$ denote the maximum size of a binary constant-weight code of word length n , minimum distance d and constant weight w . Finding the exact values of $A(n, d, w)$ (or finding lower and upper bounds for it) is a research focus in combinatorial coding theory.

Schrijver recently showed using semidefinite programming that $A(23, 8, 11) = 1288$, and the second author that $A(22, 8, 11) = 672$ and $A(22, 8, 10) = 616$. The dual solution of the corresponding semidefinite programs forces certain variables to be zero. Here we explain how this information can be used to derive uniqueness (up to code equivalence) of the codes achieving the mentioned bounds.

Gert-Jaap Polinder

Erasmus University Rotterdam
polinder@rsm.nl
Supervisor Dennis Huisman, Marie Schmidt

Title *Robust Periodic Timetabling*

Abstract In many countries, a (railway) timetable that is operated is periodic. That means that every time period, usually one hour, the timetable is repeated. The mathematical model that is used to find such a timetable is based on the Periodic Event Scheduling Problem (Serafini & Ukovich, 1989).

However, in everyday practice, disturbances occur which imply that the timetable can not be operated as planned. Many of these disturbances are only small, but when not taken into account when designing a timetable, can have a large effect. Therefore, in recent years, more research is focussed on finding a robust timetable, i.e., a timetable that is less vulnerable for disturbances. Many of the approaches so far are based on a stochastic approach or on scenario analysis. For non-periodic scheduling, more exact methods exist, but as periodic timetabling is much harder in the nominal problem already, such models have not been developed so far. Furthermore, the model size often grows when the uncertainty set becomes larger.

In this research, we wanted to fill this gap by developing a model that does not grow when the uncertainty set becomes larger. The idea of this research is based on an assignment we have done for the Robust Optimization course. We show how the model is derived and what assumptions we have. Next, we show and compare two methods how we can solve reasonable sized timetabling instances with this method.

Ernst Roos

Tilburg University

e.j.roos@uvt.nl

Supervisor Dick den Hertog

Title *A less conservative variant of Robust Optimization*

Abstract Although Robust Optimization is a great technique in dealing with uncertainty in optimization, its solutions can be too conservative when it leads to a much worse objective value than the nominal solution or even to infeasibility of the robust problem. In practice, this can lead to robust solutions being disregarded in favor of the nominal solution. This conservatism is caused by both the constraint wise approach of Robust Optimization and its core assumption that all constraints are hard for all scenarios in the uncertainty set. This paper seeks to alleviate this conservatism by proposing an alternative robust formulation that condenses all uncertainty into a single constraint that bounds the worst-case expected violation in the original constraints from above. Using recent results in distributionally robust optimization, this is shown to be a tractable formulation for both right- and left-hand side uncertainty. A computational study is performed with problems from the NETLIB library. For some problems, the uncertainty is magnified fourfold in terms of objective value of the standard robust solution, whereas we find solutions that safeguard against over half the violation while only a tenth of the uncertainty is reflected in the objective value. For problems with an infeasible standard robust counterpart, the suggested approach is still applicable and finds both solutions that safeguard against most of the uncertainty at a low price in terms of objective value.

Corné de Ruijt

VU University Amsterdam

c.a.m.de.ruijt@vu.nl

Supervisor Sandjai Bhulai

Title *A Time Dependent Job Recommendation System Using Random Survival Forest*

Abstract With the growing activity of job seekers on online professional networks, it should become easier for recruiters to find potential candidates for their open job positions. Perhaps the largest online CV database, LinkedIn, currently has over 450 million users worldwide and is still expanding. Because of this large potential, it is perhaps not surprising that the majority of organizational recruiters use social media to search and contact candidates.

As manually searching through the large number of user profiles in a CV database would be a tedious task, many methods have been proposed to (partly) automate this process, which can be seen as part of job recommendation systems. In this study we consider the specific problem of recommending candidates from a sourcing perspective. That is, we consider a recruiter who's objective is to fill an open job position by reaching out to already employed candidates listed in a large CV database, who might be interested in switching to another job.

Although most research in job recommendation systems uses semantic matching to match candidates and vacancies, we use Random Survival Forest for Competing Risk (RSF-CR) to estimate the joint probability of a candidate switching jobs within a certain amount of time, and occupying this new job for at least a certain period. The RSF is trained using an actual CV database containing over 300,00 candidates and 2 million job transitions. We validate the usefulness of RSF-CR for recommending candidates in terms of its Brier score and C-index.

Albert Schrottenboer

University of Groningen

a.h.schrottenboer@rug.nl

Supervisor Iris Vis, E. Ursavas

Title *A branch-and-price-and-cut algorithm for maintenance routing and scheduling at offshore wind farms*

Abstract We study a multi-commodity multi-period resource constrained pickup and delivery problem inspired by the short-term planning of maintenance services at offshore wind farms. To start a maintenance service, different types of scarcely available servicemen need to be delivered to the service locations. We develop resource exceeding route (RER) inequalities, whom are inspired by knapsack cover inequalities, to model the scarcity of servicemen. Besides a traditional separation approach, we present a column-dependent constraint approach to include the RER inequalities in the mathematical formulation. An alternative pricing strategy is developed to correctly include the column-dependent constraints. The resulting approach is broadly applicable for any routing problem that involves a set of scarcely available resources. We present a branch-and-price-and-cut algorithm to compare both approaches of including RER inequalities. The branch-and-price-and-cut algorithm relies on efficiently solving a new variant of the Elementary Resource Constrained Shortest Path Problem, for which a tailored pulse algorithm is developed to solve it. Computational experiments show that the RER inequalities significantly tighten the root node relaxations. The column-dependent constraint approach searches the branch and bound tree more effectively than, and appears competitive with, the traditional separation procedure. Both approaches can solve instances up to 92 nodes over 21 periods to optimality.

Jaap Storm

VU University Amsterdam

p.j.storm@vu.nl

Supervisors Michel Mandjes, Sandjai Bhulai, Wouter Kager

Title *Analysis of a Roundabout Model*

Abstract We have studied a new traffic model for a roundabout which is a hybrid between a cellular automaton and a queueing network. In particular, the model is a Markov chain. From a Markovian point of view, one might be interested in an analytical expression of the stationary distribution, and from a queueing point of view, one might be interested in performance analysis of the system. However, the circular nature of the roundabout makes the model difficult to analyze from the Markovian perspective and in turn also from the queueing perspective. We study the stability conditions of the model, i.e., when is the underlying Markov chain positive recurrent. In order to prove stability, we apply a coupling method and fluid techniques. In this talk, I will introduce the model and give an overview of the research that we did and the problems we ran into. I will also give some of the properties of the model that we have proven and post some open problems which are based on strong conjectures.

Dmitrii Usanov

CWI

usanovdml@gmail.com

Supervisor Peter van de Ven

Title *Fire Truck Relocation During Major Incidents*

Abstract In order for a fire department to effectively fight fires and respond to other incidents, it is crucial that idle fire trucks are distributed across the service area in a manner that ensures a short response time, irrespective of the location of the incident. This favourable spatial distribution of trucks may be disrupted in case of a major incident that requires many nearby fire trucks, essentially creating large gaps in the coverage. We propose an algorithm to compute the optimal relocations of idle trucks needed in order to regain good coverage. The algorithm makes relocations by solving a mathematical program that takes into account the location of the available fire trucks and the

historic spatial distribution of incidents. We apply the algorithm to the operations of the Fire Department of Amsterdam-Amstelland, and examine it against three other benchmark strategies in a simulation fitted using ten years of historical data from the fire department. We demonstrate a substantial improvement over current practice, and show that not relocating during major incidents, or relocating based on flawed heuristics and intuition, may lead to significant performance degradation.

Roland Vincze

Maastricht University

r.vincze@maastrichtuniversity.nl

Supervisor André Berger, Matthias Mnich

Title *A Polynomial-Space Algorithm for the Many-Visits Traveling Salesperson Problem*

Abstract We consider the many-visits version of the traveling salesperson problem (TSP), where, given 'n' cities, every city 'c' must be visited a certain number of times. The fastest known exact algorithm for this problem is due to Cosmadakis and Papadimitriou (SICOMP, 1984); it has time and space complexity exponential in 'n', namely $O(n^n)$. If only space polynomial in 'n' is allowed, the fastest known algorithm takes time $(n^2)^{O(n^2)}$, and is due to Grigoriev and van de Klundert (Discrete Optim., 2006).

In this paper, we improve both these results: we devise a simple algorithm using spanning tree enumeration that runs in time $O(n^n)$ and moreover requires space that is only polynomial in 'n'. We later present a way of generating all directed tree degree sequences and all directed trees for a degree sequence, further reducing the algorithm's time complexity while maintaining space complexity polynomial in 'n'.

Tom van der Zanden

Utrecht University

t.c.vanderzanden@uu.nl

Supervisor Hans Bodlaender

Title *Efficiently Computing the Shapley Value of Connectivity Games in Low-Treewidth Graphs*

Abstract At Lunteren 2017, Hebert Hamers presented work on using game-theoretic centrality measures to analyze terrorist networks. In this talk, we report on the results of a new collaboration that was set up as a result of that presentation.

Game-theoretic centrality measures are a powerful tool to identify key players in covert networks (that model, e.g., the interactions between suspected terrorists or criminals). Unfortunately, such measures are often NP-hard to compute and thus intractable, even for small graphs. However, if the graph considered has some special properties, we may be able to exploit these properties to make the computation tractable. We show that, for three connectivity games, their Shapely value can be efficiently computed if the underlying graph has low treewidth. Using this method, we are able to compute the Shapley value for several graphs for which this was previously infeasible (including, notably, the 69-vertex graph of the terrorists involved in the 9-11 attacks).

Bernard Zweers

CWI

b.g.zweers@cwi.nl

Supervisors Rob van der Mei, Sandjai Bhulai

Title *Optimizing container hinterland transportation: can we make AliExpress even cheaper?*

Abstract Containers from all over the world usually arrive in a deep-sea port, such as Rotterdam in the Netherlands. The transportation from the deep-sea port to the final destination is called hinterland transportation and contributes roughly to 50% of the total transportation. To reduce the cost of hinterland container transportation, the use of barges is getting more and more important. Barges have the advantage over trucks that they have lower costs and CO₂ - emissions per container and do not contribute to any road congestion. The latter is important, because deep-sea ports are usually located in densely populated areas. However, there is also a lot of congestion at the deep-sea terminal. Therefore, a barge cannot visit too many terminals, because otherwise it will be delayed too much and the goods will not arrive in time.

We propose a real-life operational planning problem model from a Third Party Logistic Provider (TPLP), in which the number of containers shipped per barge are maximized and the storage costs for the containers and the number of terminals visited per barge are minimized. This problem is solved with an integer linear program (ILP), yielding in strong cost reductions compared to the methods used in current practice. For real-life instances, the running time of the ILP could be too long for practical purposes, thus we have also developed a heuristic that solves the ILP in two steps: it firstly decides which terminals to visit by which barge and then assigns containers to the barges. This heuristic runs much faster than the ILP. Remarkably, for realistic real-life instances, the solution of the heuristic is almost always identical to the optimal solution of the ILP. If the heuristic solution is not optimal, it produces near optimal solutions.
