

Abstracts of the PhD presentations

Ingeborg Bikker

University of Twente

i.a.bikker@utwente.nl

Supervisor R.J. Boucherie, M. Van Houdenhoven

Title *Online capacity planning for rehabilitation treatment*

Abstract We study an online capacity planning problem in which arriving patients require a series of appointments with therapists of several disciplines, within a certain access time.

This research is motivated by a study of rehabilitation planning practices at the Sint Maartenskliniek (SMK) hospital. In practice, the prescribed treatments and activities are typically planned in the first available time slots, leaving no space for urgent patients who require a series of appointments at short notice. This leads to cancellations of appointments or long access times for urgent patients, which have a negative effect on the quality of care and on patient satisfaction.

The objective of our research is to plan capacity for the appointment series of a patient at the moment of his or her arrival, in such a way that the total number of requests planned within their required access time is maximized. We formulate this problem as a Markov decision process that takes into account the current state of already planned capacity, and predicted future arrivals. An approximate dynamic programming approach is used to obtain approximate solutions.

Nardo Borgman

University of Twente

n.j.borgman@utwente.nl

Supervisor Erwin Hans, Richard Boucherie

Title *Appointment scheduling with unscheduled arrivals and reprioritization*

Abstract Appointment scheduling is a well studied problem in the literature. Inspired by the real life problem of a radiology department in a Dutch hospital, we study the problem of scheduling appointments, taking into account unscheduled arrivals and reprioritization. The radiology department offers CT diagnostics to both scheduled and unscheduled patients. Of these unscheduled patients, some must be seen immediately, while others may wait for some time. Herein a trade-off is sought between acceptable waiting times for appointment patients and unscheduled patients' lateness. In this paper we propose a model to calculate the performance of a given appointment schedule in terms of waiting time and lateness. Also we propose a constructive and local search heuristic that embeds this model and optimizes the schedule. In addition we present computational results using instances from the case study in the Dutch hospital. These results show a considerable decrease of waiting time is possible for scheduled patients, while still treating unscheduled patients on time.

Sem van Brummelen

University of Twente

s.p.j.vanbrummelen@utwente.nl

Supervisor Nico van Dijk

Title *Time dependent waiting time computation at Dutch blood collection sites.*

Abstract The Dutch blood bank, Sanquin, collects blood from donors on a non-remunerated basis. It is therefore important to keep waiting times as short as possible. Standard analytic models to calculate waiting times, use some form of steady state distribution. If arrivals are nicely spread out

over the day, this would be reasonably justified. But, as Dutch blood collection sites allow free walk ins, the arrival process of donors is highly time dependent. Peaks in arrival patterns show up early in the morning, around noon and just before dinner time. Accurately approximating waiting times is especially important during these peak times. A transient rather than steady state approach is therefore required.

Anne Buijsrogge

University of Twente

a.buijsrogge@utwente.nl

Supervisor Pieter-Tjerk de Boer, Werner Scheinhardt

Title *On the state-independent change of measure for the G|G|1 tandem queue*

Abstract Rare event simulation has been of interest for more than two decades and is, for example, of interest in telecommunication. In importance sampling the rare event is made less rare by changing the underlying probability distribution. This is called a change of measure. In this talk I will discuss the state-independent change of measure for the G|G|1 tandem queue as proposed by Parkeh and Walrand in 1989. In the G|G|1 tandem queue, the rare event of our interest is the probability to reach the level N in a busy cycle. For the M|M|1 tandem queue it has been shown previously that the change of measure proposed by Parekh and Walrand is not necessarily asymptotically efficient. By considering specific sample paths, we will provide necessary conditions for a change of measure for the G|G|1 tandem queue to be asymptotically efficient.

Ewan Cahen

CWI

ewan.cahen@cwi.nl

Supervisor Bert Zwart, Michel Mandjes

Title *Rare event analysis and efficient simulation for a multi-dimensional ruin problem*

Abstract We look at large deviations of multivariate stochastic processes, in particular, we consider the event that both components of a bivariate stochastic process ever simultaneously exceed some large level; a leading example is that of two Markov fluid queues driven by the same background process ever reaching some large level u . Exact analysis being prohibitive, we resort to asymptotic techniques and efficient simulation. The first result we present concerns various expressions for the decay rate of the probability of interest, which are valid under G\''artner-Ellis-type conditions. The second result is an importance-sampling-based rare-event simulation technique for the bivariate Markov modulated fluid model, which is capable of asymptotically efficiently estimating the probability of interest. We illustrate this with a number of numerical experiments.

Fabio Cecchi

Eindhoven University of Technology

F.Cecchi@tue.nl

Supervisor Sem Borst, Johan van Leeuwen

Title *CSMA Networks in a Many-Sources Regime: A Mean-Field Approach*

Abstract With the rapid advance of the Internet of Everything, both the number of devices and the range of applications that rely on wireless connectivity show huge growth.

Driven by these pervasive trends, wireless networks grow in size and complexity, supporting immense numbers of nodes and data volumes, with highly diverse traffic profiles and performance requirements. While well-established methods are available for evaluating the throughput of

persistent sessions with saturated buffers, these provide no insight in the delay performance of flows with intermittent packet arrivals. The occurrence of empty buffers in the latter scenario results in a complex interaction between activity states and packet queues, which severely complicates the performance analysis. Motivated by these challenges, we develop a mean-field approach to analyze buffer contents and packet delays in wireless networks in a many-sources regime.

The mean-field behavior simplifies the analysis of a large-scale network with packet arrivals and buffer dynamics to a low-dimensional fixed-point calculation for a network with saturated buffers. In particular, the analysis yields explicit expressions for the buffer content and packet delay distribution in terms of the fixed-point solution.

Extensive simulation experiments demonstrate that these expressions provide highly accurate approximations, even for a fairly moderate number of sources.

Kevin Dalmeijer

Erasmus University Rotterdam

dalmeijer@ese.eur.nl

Supervisor Albert P.M. Wagelmans, Remy Spliet

Title *A branch-and-cut algorithm for the Time Window Assignment Vehicle Routing Problem*

Abstract The Vehicle Routing Problem with Time Windows (VRPTW) is the problem of routing vehicles such that deliveries are made at minimum cost, taking predefined time windows and vehicle capacity into account. We consider a generalization of the VRPTW in which customer demand is no longer deterministic, but randomly drawn from a finite set of scenarios. Furthermore, we allow changing the (endogenous) time windows, but only *before* the demand scenario becomes known. Each endogenous time window is required to be within the client specific exogenous time window. This generalization is called the Time Window Assignment Vehicle Routing Problem (TWAVRP) and is recently introduced in the literature.

To the best of our knowledge, the TWAVRP has only been solved to optimality by formulating it as a set-partitioning type problem and solving it by a branch-price-and-cut algorithm. Using this algorithm, instances up to 25 customers and three demand scenarios can be solved within one hour of computation time. In this paper we use a different formulation and solve the TWAVRP using a branch-and-cut algorithm. Numerical experiments show that for all test-instances the new algorithm is superior to the branch-price-and-cut algorithm in the literature, and an average speedup factor of over 180 is achieved. We show that a significant part of this speedup can be attributed to the use of precedence inequalities, novel valid inequalities that we introduce specifically for the TWAVRP. Tests on larger instances show that with the new algorithm, instances up to 35 customers and three demand scenarios can be solved to optimality within one hour of computation time.

Bas Dietzenbacher

Tilburg University

B.J.Dietzenbacher@uvt.nl

Supervisor P.E.M. Borm

Title *Decomposition of Network Communication Games*

Abstract Using network control structures this paper introduces network communication games as a generalization of vertex games and edge games corresponding to communication situations and studies their decomposition into unanimity games. We obtain a relation between the dividends of the network communication game and the underlying transferable utility game, which depends on the structure of the undirected graph. This relation extends the computational results for cycle-free networks to general undirected graphs and is used to derive new characterizations of the Myerson

value and the position value. Moreover, network communication games also allow to consider the vertices and the edges of the graph as players simultaneously, leading to a new network value.

Annette Ficker

KU Leuven

annette.ficker@kuleuven.be

Supervisor Frits Spieksma

Title *Balanced Optimization with Vector Costs*

Abstract We consider so-called balanced optimization problems with vector costs. We propose a framework containing such problems; this framework allows us to investigate the complexity and approximability of these problems in a general setting. More concrete, each problem in the framework admits a 2-approximation, and for many problems within the framework this result is best-possible, in the sense that having a polynomial-time algorithm with a performance ratio better than 2 would imply $P=NP$. Special attention is paid to the balanced assignment problem with vector costs.

Stijn Fleuren

Eindhoven University of Technology

S.T.G.Fleuren@tue.nl

Supervisor Erjen Lefeber, Ivo Adan

Title *Optimizing pre-time traffic light control and lane-use-arrows at isolated intersections*

Abstract When a traffic intersection is designed, this design is usually based on the forecasted demands at the intersection in the upcoming years. Typically the current demands are estimated or counted and some growth factor is applied to estimate the future demands. An important problem in the design of a traffic intersection is deciding what lane-use arrows to use for each lane at the intersection. These lane-use-arrows indicate what movements are allowed at that specific lane, e.g., are you allowed to make a right turn? It is desirable to choose these lane-use-arrows so that the maximum sustainable growth factor is as large as possible; the time period during which the intersection can handle the amount of traffic that arrives at it is then as large as possible. To this end, a mixed-integer programming problem is formulated using the cycle periodicity formulation of Nachtigall [1]. This optimization problem chooses the lane-use-arrows such that the growth factor is maximized. Besides optimizing the lane-use-arrows it also optimizes the pre-time traffic control, i.e., it returns a signal group schedule. Such a signal group schedule visualizes when each of the traffic-lights have a green, yellow and red indication. The optimization returns a signal group schedule that can handle the forecasted demands for the maximized growth factor.

[1] K. Nachtigall. A branch and cut approach for periodic network programming. Inst. für Mathematik, 1994.

Wouter van Heeswijk

University of Twente

w.j.a.vanheeswijk@utwente.nl

Supervisor Martijn Mes, Marco Schutten

Title *An Approximate Dynamic Programming Algorithm for the Delivery Dispatching Problem*

Abstract We address the dispatch decision faced by an urban consolidation center. The center receives orders according to a stochastic order arrival process, and dispatches orders for the last-mile

distribution in batches. The operator of the center aims to find the cost-minimizing consolidation policy, depending on inventory at hand, pre-announced orders, and stochastic arrivals. We formulate this problem as a variant of the Delivery Dispatching Problem that includes dispatch windows, and formulate it as a Markov model. For toy-sized instances, we solve this model to optimality. Larger instances have an intractably large state space, outcome space, and action space. We describe an Approximate Dynamic Programming (ADP) algorithm that can handle such instances. Value function approximation is used to estimate the downstream costs. To handle large action spaces, we formulate an integer linear program to be used within our ADP algorithm. Through numerical experiments, we show for small instances that we closely approximate the optimal values. To evaluate the performance of our ADP policies for larger instances, we test against various benchmark policies, including a policy based on scenario sampling. Particularly when the dispatch problem offers sufficient flexibility in dispatch times, ADP clearly outperforms the benchmark policies.

Pim van der Hoorn

University of Twente

w.l.f.vanderhoorn@utwente.nl

Supervisor Nelly Litvak

Title *Phase transition for average number of erased edges in directed configuration model*

Abstract Models for generating directed simple graphs are important for the analysis of real-world networks. One such model is the directed erased configuration model. Here each node is first assigned a number of outgoing and incoming stubs. Then outgoing stubs are paired with a incoming stub, selected uniformly at random from among all available stubs to create an edge. Afterwards, all self loops are removed and all multiple edges between nodes are merged.

It is clear that this removal step alters the degrees, which can give rise to finite size effects. Still, when stubs are sampled according to an i.i.d. procedure from some given distributions, it is well known that the empirical degree distributions of the resulting graph converge to the original ones as the size of the graph goes to infinity. However, the exact speed of this convergence, or the finite size effect of the removal of edges on other processes has remain unknown.

In this talk I give a first characterization of the speed of convergence, by providing upper bounds for the average number of erased edges, in terms of the size of the graph. Remarkably, when the degree distributions follow a power-law, we observe three different scaling regimes, depending on the exponents of the distributions. These results provide a first crucial step towards evaluation of finite size effects in networks.

Asparuh Hristov

CWI

A.V.Hristov@cwi.nl

Supervisor Rob van der Mei, Sandjai Bhulai

Title *Throughput and Bottleneck Analysis of Tandem Queues with Nested Sessions*

Abstract Various types of systems across a broad range of disciplines contain tandem queues with nested sessions. Strong dependence between the servers has proved to make such networks complicated and difficult to study. Exact analysis is in most of the cases intractable. Moreover, even when one is given performance metrics such as the saturation throughput and the utilization rates of the servers, determining the limiting factor of such a network can be far from trivial. In our work, we present a simple, tractable and nevertheless relatively accurate method for approximating the above mentioned performance measurements for any server in a given network. In addition, we show how one can use those values in order to derive a novel metric, which can be used for bottleneck identification.

David Koops

University of Amsterdam

koops.david@gmail.com

Supervisor Michel Mandjes, Onno Boxma

Title *Lévy Tandem Networks in Heavy-Traffic Regimes*

Abstract We study the joint stationary distribution of the workloads in a fluid tandem queue using heavy-traffic analysis. We consider different types of input, ranging from compound Poisson to α -stable input, with $1 < \alpha < 2$, and Brownian input. Similar to known single server heavy-traffic results, we find a dichotomy between finite and infinite variance input processes.

A special feature is that we do not only consider the usual heavy-traffic regime, in which the load goes to unity for a particular server, but also a regime in which we increase the load of both servers to one. It appears that this leads to different heavy-traffic asymptotics. In a special case it is shown by a simulation study that the simultaneous heavy-traffic approximation performs better than the usual heavy-traffic approximation.

Peter Kovacs

University of Amsterdam

P.Kovacs@uva.nl

Supervisor Rudesindo Nunez Queija

Title *Routing policies for a partially observable two-server queueing system*

Abstract We consider a queueing system controlled by decisions based on partial information. The motivation for this work stems from road traffic, in which drivers may, or may not, be subscribed to a smartphone application for dynamic route planning.

Our model consists of two queues, both operating in a FIFO manner with independent exponential service times, serving two types of jobs. Arrivals occur according to a Poisson process of rate λ ; a fraction α of jobs (type X) are observable and controllable. At all times the number of X jobs in each queue and their individual positions are known. Upon its arrival a router decides which queue the next X job should join. Y jobs (fraction $1-\alpha$) are non-observable, not even the number of Y jobs in the queues, and non-controllable. They join queue i with static routing probabilities $p^{\alpha Y}_i$.

We address the following research questions: 1) what penetration level α (percentage of X jobs) is needed for effective control, 2) which policy should be implemented at the router, and 3) what is the added value of having more system information (e.g., average service times and the values of $p^{\alpha Y}_i$)?

We investigate these questions through extensive simulations. This simulation study reveals that for heavily loaded systems a low penetration level suffices and that the performance of a simple policy that relies on little system information is close to the optimal w -JSQ (weighted join-the-shortest-queue policy). The latter result is confirmed by analysis of a dynamical system that approximates the stochastic evolution in heavy traffic.

Gréanne Leefink

University of Twente

a.g.leefink@utwente.nl

Supervisor Erwin Hans, Richard Boucherie, Ingrid Vliegen

Title *Histopathology laboratory operations analysis and improvement*

Abstract Histopathology laboratories aim to deliver high quality diagnoses based on patient tissue samples. Indicators for quality are the accuracy of the diagnoses and the diagnostic turnaround

times. However, challenges exist regarding employee workload and turnaround times, which both impact the diagnostic quality. We propose a decomposed planning and scheduling method for the histopathology laboratory using (mixed) integer linear programming to improve the spread of workload and reduce the diagnostic turnaround times. First, the batching problem is considered, in which batch completion times are equally divided over the day to spread the workload. This reduces the peaks of physical work available in the laboratory. Thereafter, the remaining processes are scheduled to minimize the tardiness of orders. Results show that using this decomposed method, the peaks in histopathology workload in UMC Utrecht, a large university medical center in the Netherlands, are reduced with up to 50% by better spreading the workload over the day. Furthermore, turnaround times are reduced with up to 20% compared to current practices. This approach is currently being implemented in the histopathology laboratory of UMC Utrecht.

Daphne van Leeuwen

CWI

dleeuwe@cwi.nl

Supervisor Rob van der Mei, Sindo Nunez Queija, Sandjai Bhulai

Title *Near-optimal switching strategies for a tandem queue*

Abstract Motivated by various applications in logistics, road traffic and production management, we investigate two versions of a tandem queueing model in which the service rate of the first queue can be controlled. The objective is to keep the mean number of jobs in the second queue as low as possible, without compromising the total system delay (i.e. avoiding starvation of the second queue). The balance between these objectives is governed by a linear cost function of the queue lengths. In the first model, the server in the first queue can be either switched on or off, depending on the queue lengths of both queues. This model has been studied extensively in the literature. Obtaining the optimal control is known to be computationally intensive and time consuming. We are particularly interested in the scenario that the first queue can operate at larger service speed than the second queue. This scenario has received less attention in literature. We propose an approximation using an efficient mathematical analysis of a near-optimal threshold policy based on a matrix-geometric solution of the stationary probabilities that enables us to compute the relevant stationary measures more efficiently and determine an optimal choice for the threshold value. In some of our target applications, it is more realistic to see the first queue as a (controllable) batch-server system. We follow a similar approach as for the first model and obtain the structure of the optimal policy as well as an efficiently computable near-optimal threshold policy. We illustrate the appropriateness of our approximations using simulations of both models.

Siqiao Li

VU University Amsterdam

aprilsiqiao522@gmail.com

Supervisor Ger Koole

Title *Different Control Policies of Flexible Capacity Allocation for Multi-type Patients in Radiotherapy Process*

Abstract The re-entrance of cancer patients in radiotherapy treatment phase makes the process specific and challenging to operation research study. This paper deals with capacity allocation of the most important resource -- Linear Accelerators (LINACs) for multi-type patients. Suggested by hospital managers, we first divided patients into two groups according to the waiting time target (urgent/ routine), then a static capacity allocation scheme for a given service level is given. For better use of the capacity, we introduced flexible capacity allocation schemes, which allowed two patient groups shared the capacity with some threshold control policy. In the paper, different control policies

and also some classic static priority scheduling policies are compared considering the performance measurements including the total amount of capacity needed to reach the service levels, average waiting time and breaching probability for each patient type and the robustness of the policy. Our model is based on the 'slot server perspective', which considered the time slot on the LINACs as server and patients' re-entrance times can be considered as service time. With the fact that there are multi-types of patients with various waiting time targets and treatment protocols in each patient group, the process can be considered as patient dispatching problem in M/G/s queueing model, leading to no analytical results such as the probability of waiting time and even average waiting time. Instead, we use simulation and heuristic algorithm to analyze the system and derive optimal threshold value for each control policy. The results showed that some control policies outperform the static priority policies in different aspects, which is very delightful since these simple policies are more convenient to use in the reality.

Britt Mathijsen

Eindhoven University of Technology

b.w.j.mathijsen@tue.nl

Supervisor Bert Zwart, Johan van Leeuwen

Title *Capacity allocation in a transient queue.*

Abstract Performance analysis of queueing systems is typically done through its steady-state distribution and its derivatives. Accordingly, staffing rules rely on these measures. However, in practice system parameters are rarely constant over time nor do processes run long enough to reach equilibrium. As a consequence, allocating capacity based on stationary measures could result in performance discrepancies. We study a cost minimization problem of a single-server queue with Levy input, over a finite time horizon, starting out of equilibrium. By quantifying the impact of the transient phase of the process, we come up with a correction to the objective function. Subsequently, we show how this novel approximation results in a staffing rule which is a refinement of the optimal value in stationarity. Finally, we analyze the associated optimality gap and determine in which settings our correction is necessary.

Thomas Meyfroyt

Eindhoven University of Technology

t.m.m.meyfroyt@tue.nl

Supervisor Sem Borst, Onno Boxma

Title *Analysis of Degree-Dependent Counter-Based Gossip Protocols on Treelike Networks*

Abstract In wireless networks, broadcasting is a basic functionality used to support many applications, like neighbor discovery, data dissemination, data aggregation and more. To this end, gossip protocols, have emerged as a reliable approach to implement broadcast. In particular, the use of counter-based gossiping schemes has proven to be a powerful technique for implementing highly scalable and robust services. Such schemes suppress broadcasts by a node, if the number of broadcasts it has already received exceeds some predefined threshold. In this presentation we analyze the probability that a node's broadcast is suppressed during a single broadcasting round when a degree-dependent counter-based gossiping scheme is used. Specifically, we analyze the suppression probability of nodes in an infinite tree network, given the degree distribution of the tree. Furthermore, we propose an algorithm that determines how to configure the degree-dependent thresholds on infinite tree networks in order to reach some desired target suppression probabilities. Lastly, we show that the thresholds generated by the algorithm also perform remarkably well on finite non-tree networks and are robust to changes in the node-degree distribution and network topology.

Frans de Ruiter

Tilburg University

F.J.C.T.deRuiter@uvt.nl

Supervisor Dick den Hertog

Title *Duality in Two-stage Adaptive Linear Optimization: Faster Computation and Stronger Bounds*

Abstract We derive and exploit duality in general two-stage adaptive linear optimization models. The equivalent dualized formulation we derive is again a two-stage adaptive linear optimization model. Therefore, all existing solution approaches for two-stage adaptive models can be used to solve or approximate the dual formulation. The new dualized model differs from the primal formulation in its dimension and uses a different description of the uncertainty set. We show that the optimal primal affine policy can be directly obtained from the optimal affine policy in the dual formulation. We provide empirical evidence that the dualized model in the context of two-stage lot-sizing on a network and two-stage facility location problems solves an order of magnitude faster than the primal formulation with affine policies. We also provide an explanation and associated empirical evidence that offer insight on which characteristics of the dualized formulation make computations faster. Furthermore, the affine policy of the dual formulations can be used to provide stronger lower bounds on the optimality of affine policies.

Jori Selen

Eindhoven University of Technology

j.selen@tue.nl

Supervisor Ivo Adan, Johan van Leeuwen

Title *Approximate performance analysis of generalized join the shortest queue routing*

Abstract We propose a highly accurate approximate performance analysis of a heterogeneous server system with a processor sharing service discipline and a general job-size distribution under a generalized join the shortest queue (GJSQ) routing protocol. The GJSQ routing protocol is a natural extension of the well-known join the shortest queue routing policy that takes into account the non-identical service rates in addition to the number of jobs at each server. The performance metrics that are of interest here are the equilibrium distribution and the mean and standard deviation of the number of jobs at each server. We show that the latter metrics are near-insensitive to the job-size distribution using simulation experiments. By applying a single queue approximation we model each server as a single server queue with a state-dependent arrival process, independent of other servers in the system, and derive the distribution of the number of jobs at the server. These state-dependent arrival rates are intended to capture the inherent correlation between servers in the original system and behave in a rather atypical way.

Matteo Seminaroti

CWI

M.Seminaroti@cwi.nl

Supervisor Monique Laurent

Title *A new simple recognition algorithm for Robinsonian matrices*

Abstract In many applications it is required to order a given set of objects when the only available information among the objects is their pairwise similarity (or dissimilarity). Robinsonian matrices arise in the classical seriation problem and play an important role in many applications where unsorted similarity (or dissimilarity) information must be reordered. A Robinson similarity matrix is a

symmetric matrix whose entries increase monotonically along rows and columns when moving towards the diagonal. A Robinsonian similarity matrix is then a symmetric matrix whose rows and columns can be permuted obtaining a Robinson similarity matrix.

We present a new simple recognition algorithm for Robinsonian matrices. Given the equivalence between the class of unit interval graphs and 0-1 Robinsonian matrices, we extend the famous 3-sweep algorithm of Corneil for recognizing unit interval graphs. We introduce a generalization of Lexicographic Breadth-First Search (Lex-BFS) for weighted graphs, named Similarity First-Search (SFS), and we present a multisweep polynomial time recognition algorithm based on a series of consecutive SFS sweeps. The algorithm is extremely simple and, given a symmetric matrix of size n with m nonzero entries, we show that it returns a Robinson ordering of the matrix after at most n sweeps.

Berkas Serbetci

University of Twente

b.serbetci@utwente.nl

Supervisor J. Goseling, R. J. Boucherie

Title *On Optimal Geographical Caching in Heterogeneous Cellular Networks*

Abstract There is extreme growth in data traffic over cellular networks and the growth rate of the demand will increase in the upcoming years such that current and near-future cellular network architectures will not be able to support it. A solution to this problem is to replace backhaul capacity with the storage capacity at small cell access points. Thus, part of the data is stored at the wireless edge and the backhaul is used only to refresh this partly stored data.

In this work, we investigate optimal geographical caching in heterogeneous cellular networks. There is a finite available content containing files with the same size and with different popularities following the Zipf distribution. There are different types of base stations with different cache size capacities in the plane. The performance metric is the total hit probability which is the probability that the user of interest will find the content that he requires to gather in one of the base stations that he is covered by. The user of interest is located at an arbitrary location in the plane and may be covered by a set of different types of base stations or may not be covered at all. We consider various scenarios for the distribution of base stations in the plane. We show that the optimal content placement problem for a new type of base station, when the other types of base stations use a pre-determined content placement policy, is a convex optimization problem. We analyze what to cache in the new type of base station and show how the hit probability evolves as the deployment density of the new type of base station varies.

Nicos J. Starreveld

University of Amsterdam

N.J.Starreveld@uva.nl

Supervisor M.M. Mandjes, R. Bekker

Title *Occupation times for stochastic processes and applications*

Abstract In this paper we consider a stochastic process $\{X(t) : t \geq 0\}$ taking values on the state space (E, \mathcal{E}) and a partition of the state space $E = A \cup B$ into two disjoint regions A, B . We are interested in the occupation time, denoted by $\alpha(t)$, of state A up to time t , denoted by

$$\alpha(t) = \int_0^t \mathbf{1}_{\{X(s) \in A\}} ds.$$

We analyze both the transient and the asymptotic behavior of the occupation time $\alpha(t)$, the transient behavior by studying its distribution function and its double transform. For the asymptotic behavior we prove a central limit theorem and a large deviations principle.

The motivating examples stem from queueing theory, storage problems and finance. We discuss in detail the case that the process $X(t)$ is a spectrally positive Levy process and a spectrally positive Levy process reflected at its latest infimum. For the first case we establish a distributional equality between the occupation time of the negative half plane and the first epoch at which the supremum is attained.

Veerle Timmermans

Maastricht University

v.timmermans@maastrichtuniversity.nl

Supervisor Tobias Harks

Title *Uniqueness of Equilibria in Atomic Splittable Polymatroid Congestion Games*

Abstract We revisit the issue of uniqueness of equilibria in atomic splittable congestion games. In this class of games there is a finite set of resources and a finite set of players, and each player is associated with a weight and a collection of allowable subsets of resources. A strategy for a player is a (possibly fractional) distribution of the weight over the allowable subsets.

We derive a uniqueness result based on polymatroid theory: when the strategy space of every player is a bidirectional flow polymatroid, then equilibria are unique. Bidirectional flow polymatroids are introduced as a subclass of polymatroids possessing certain exchange properties. We show that important cases such as base orderable matroids, and therefore gammoids, transversal matroids and graphic matroids on generalized series-parallel graphs, can be recovered as a special case of bidirectional flow polymatroids.

On the other hand we show that matroidal set systems are in some sense necessary to guarantee uniqueness of equilibria: for every atomic splittable congestion game with at least three players and non-matroidal set systems per player, there is an isomorphic game having multiple equilibria. Our results leave a gap between base orderable matroids and general matroids for which we do not know whether equilibria are unique.

Denise Tönissen

Eindhoven University of Technology

D.D.Tonissen@tue.nl

Supervisor Geert-Jan van Houtum, Joachim Arts

Title *Maintenance Location Routing for Train Units*

Abstract The Maintenance Location Routing Problem (MLRP) for Train Units is a problem where we locate maintenance locations, while also taking the maintenance routing into account. To our best knowledge, this is a new problem which has never been thoroughly studied before. We argue that making facility location and maintenance routing decisions jointly is important for this problem and study approaches for similar problems in the literature. We identify several complicating features of a joint approach in the railway environment. The most important of these is that routing feasibility is more of an issue than minimizing transportation cost.

We suggest methods that deal with all these features simultaneously. The first method uses the following steps. First, generate feasible routes to possible candidate facility locations by solving a multi-commodity flow problem for every train type in a rolling horizon fashion. Second, open candidate locations by combining feasible location-routes for all train types. Such combinations can be found by solving a generalization of the hitting set problem. The second method simplifies the maintenance routing by assuming there are only two ways to reach the maintenance facility. First, the

train unit reaches the maintenance facility with successful maintenance routing without any additional costs. Second, the train unit reaches the facility by deadlocking, with as cost the normal transportation costs and additional cost associated with the unavailability of the train unit. In a preprocessing step, in which we consider all possible maintenance routing paths, we calculate the occurrence probabilities and shortest route deadlock costs of the second option. This allows us to calculate the expected maintenance routing costs per train unit to every candidate facility, which can be used as the ‘transportation’ costs in a standard uncapacitated facility location model. We also discuss possible extensions of the problem such as dealing with unexpected failures, capacity sizing for maintenance locations and a joint problem where the maintenance routing possibilities are not given, but can be actively improved.

Konstantin Vandyshev

Delft University of Technology

K.Vandyshev@tudelft.nl

Supervisor Karen Aardal, Dion Gijswijt

Title *Phase Shifter Transformer Optimization within Flow Based Market Coupling Methodology*

Abstract The talk is devoted to an optimization procedure which uses Phase Shifter Transformers (PST) within the scope of the Flow Based (FB) Market Coupling (MC) methodology in the power system management. The goal of optimization is to increase the available capacity for energy exchange per border in a certain direction. The advantage of the FB MC environment is its ability to easily identify the most relevant branch flow constraints, that limit the transportation of energy from one country (or region) to another. The analysis takes into account the information on the critical branches that might be overloaded, and therefore the obtained solution is secure. The main conclusion of the work is that the proper usage of PST may increase available transfer capacities, thus allowing more energy to be transported in the desired border direction. This procedure can be applied as an operational tool to find the best system settings for a particular purpose. Alternatively, it can be used for comparison of investment proposals.

Marjolein Veenstra

University of Groningen

marjolein.veenstra@rug.nl

Supervisor Iris F. A. Vis, Kees Jan Roodbergen

Title *The Pickup and Delivery Traveling Salesman Problem with Handling Costs*

Abstract This paper introduces the pickup and delivery traveling salesman problem with handling costs (PDTSPH). In the PDTSPH, a single vehicle has to transport loads from origins to destinations. Loading and unloading of the vehicle is operated in a last-in-first-out (LIFO) fashion. However, if a load must be unloaded that was not loaded last, additional handling operations are allowed to unload and reload other loads that block access. Since the additional handling operations take time and effort, penalty costs are associated with them. The aim of the PDTSPH is to find a feasible route such that the total costs, consisting of travel costs and penalty costs, are minimized. We show that the PDTSPH is a generalization of the pickup and delivery traveling salesman problem (PDTSP) and the pickup and delivery traveling salesman problem with LIFO loading (PDTSPL). We propose a large neighborhood search (LNS) heuristic to solve the problem. We compare our LNS heuristic against best known solutions on 163 benchmark instances for the PDTSP and 42 benchmark instances for the PDTSPL. We provide new best known solutions on 52 instances for the PDTSP and on 14 instances for the PDTSPL, besides finding the optimal or best known solution on 100 instances for the PDTSP and on 21 instances for the PDTSPL. The LNS finds optimal or near-optimal solutions on instances for the

PDTSPH. Results show that PDTSPH solutions provide large reductions in handling compared to PDTSP solutions, while increasing the travel distance by only a small percentage.

Andrej Winokurow

Maastricht University

a.winokurow@maastrichtuniversity.nl

Supervisor Alexander Grigoriev, André Berger

Title *Location, pricing and the problem of Apollonius*

Abstract In Euclidean plane geometry, Apollonius' problem is to construct a circle in a plane that is tangent to three given circles. We will use a solution to this ancient problem to solve several versions of the following geometric optimization problem. Given is a set of customers located in the plane, each having a demand for a product and a budget. The task is to determine location of production facilities in the plane and one price for the product, such that the revenue generated from those customers, for which the sum of travel and purchase costs does not exceed their budget, is maximized.

Firstly, we address the location-pricing problem with one facility and three customers having unit demands. It is solved by using a solution to the Apollonius' problem. Secondly, we extend the algorithm to the case with n customers having arbitrary demands and prove that it can be solved in $O(n^3)$ time. Thirdly, we present an exact algorithm solving the location-pricing problem with m facilities in $O(n^{3m+1})$ time. Finally, we consider the non-uniform location-pricing problem.

Qing Chuan Ye

Erasmus University Rotterdam

ye@ese.eur.nl

Supervisor Rommert Dekker

Title *Fair task allocation in transportation*

Abstract Task allocation problems have traditionally focused on cost optimization. However, more and more attention is being given to cases in which cost should not always be the sole or major consideration.

In this paper, we study a fair task allocation problem in transportation where an optimal allocation not only has low cost but more importantly, it distributes tasks as even as possible among heterogeneous participants who have different capacities and costs to execute tasks. To tackle this max-lexmin fair minimum cost allocation problem (MFMCA), we analyze and solve it in two parts using two novel polynomial-time algorithms. We show that despite the new fairness criterion, the proposed algorithms can solve the MFMCA problem optimally in polynomial time.

In addition, we conduct an extensive set of experiments to investigate the trade-off between cost minimization and fairness. Our experimental results demonstrate the benefit of factoring fairness into task allocation. Among majority of the test instances, fairness comes with a very small price in terms of cost.

Sha Zhu

Erasmus University Rotterdam

zhu@ese.eur.nl

Supervisor Rommert Dekker

Title *Using Maintenance Plan in Spare Part Demand Forecasting and Inventory Control*

Abstract Lumpiness in spare parts demand, sometimes can be extremely heavy, is a nightmare for

inventory managers since it increases the difficulty in demand forecasting largely. In the situation of preventive maintenance, the lumpiness in spare part demand is triggered by both of the lumpiness in component repairs and the uncertainty of individual probable defective component generating spare part demand. Since maintenance plan provides prior information about defective component arrival, it captures the lumpiness in component repairs and further explains the lumpiness in spare parts demand to some extent. In this way, maintenance plan might prevent us from keeping redundant inventory in the period when there is no component replacement and allow us to estimate the spare part demand based on component arrivals. Therefore, we can make good use of maintenance plan to improve spare part demand forecasting and inventory management. It is valuable especially for critical spare parts which are characterized as expensive, and slow-moving items with high stockout costs and a high risk of obsolescence.

We propose an estimation of spare part demand distribution and build a dynamic model to obtain the order policy. We first propose a periodic review, lost sale inventory model which minimizes total inventory holding, shortage penalty and ordering cost under maintenance plan. Next, we explore the relationship between the component repairs and the spare part demand. Assuming the same spare part installed in the same type of component has a constant replacement probability in each time period, we can estimate the probability that a component repair needs a certain spare part. With the information of component arrivals by maintenance plan, the number of spare parts demand is then binomial distributed and we can obtain the order policy from the inventory model. We consider two solution concepts, one with and another without maintenance plan, in our inventory model to examine the value of information. A case study performed on real data provided by Fokker Service show that the total cost is reduced by 8.4% using maintenance plan.
