

Titles and Abstracts

1. Danyu Bai, Dalian Maritime University, China

Title: Flow Shop Learning Effect Scheduling Problem with Release Dates

Abstract: In a real-world assembly environment, the components of a product arrive at a plant over time. Works-in-process are assembled into end-products by following an identical processing route. When a worker at a particular stage repeatedly handles similar tasks and gains the knowledge to execute a task efficiently, the processing time for later tasks is shortened significantly. This assembly process can be described as the flow shop learning effect scheduling problem with release dates, in which the learning effect is dependent on position. The objective is to minimize one of three different criteria, namely, makespan, total completion time and total quadratic completion time. This scheduling problem is formulated as a mixed integer programming (MIP) model. For small-scale problems, a branch and bound (B&B) algorithm with an efficient branching rule is proposed to obtain optimal solutions. The MIP model and the B&B algorithm provide key evidence for academic research. For large-scale problems, the asymptotic optimality of a class of shortest processing time available (SPTA)-based heuristics is proven in terms of probability limit. The convergence property indicates that an SPTA-based heuristic can serve as an optimal schedule under the industrial setting, where thousands of tasks are typically executed on a set of machines.

2. Ruibin Bai, University of Nottingham Ningbo, China

Title: Lagrange Dual Bound Computation for Stochastic Service Network Design: the Advantage of Scenario Bundling

Abstract: We consider the Lagrange dual problem of stochastic service network design (SSND) formed by relaxing the non-anticipativity constraints for obtaining lower bounds. The recently proposed FW-PH, which combines progressive hedging with a Frank-Wolfe-like method, is a theoretically appealing algorithm to solve the Lagrange dual problem. Although convergence to the optimal Lagrange dual value is assured, the number of iterations and the run-time required by the FW-PH could still be substantial. In order to obtain a tighter lower bound within less run-time, we extend the FW-PH by replacing scenario-wise decomposition with bundle-wise decomposition, which divides the problem into multiple scenario-bundle subproblems, rather than by individual scenarios into single-scenario subproblems. Our theoretical analysis shows that this extension not only keeps true the convergence property, but gives tighter La-

grange dual bound at convergence. Computational experiments on the test instances for SSND demonstrate that the extended FW-PH is far superior to the basic version, providing either a dramatic saving in the run-time required to achieve convergence or a much tighter lower bound when terminated due to the time limit.

3. Jacek Blazewicz, Poznan University of Technology, Poland

Title: Cloud Brokering: Grouping Cloud Services into Packages

Abstract: Cloud computing quickly gains great interest and acceptance of both practitioners and the scientific community. Thanks to this paradigm, on-demand access to distributed resources and storage space can be realized in a transparent and smooth manner.

With the appropriate implementation, the use of a wide range of services is possible without involvement in the hardware or system configuration. Emerging new business models offer free services enabling the use of data (coming from system users) for marketing purposes. The number of offered services is growing rapidly. There will be also more and more service providers who compete for the market and offer similar services. Small differences cause some confusion and the best choice is a growing challenge for the customer. On the other hand, it should be emphasized that the multitude of available services gives cloud service providers / sellers the possibility to create solid IT solutions based on security, reliability and trust, while minimizing total costs.

It is becoming increasingly difficult for end-users to choose the right services at the best price. Indeed, numerous cloud services look-alike but in reality they differ in terms of price, performance or reliability. One of the answers to this problem is cloud service brokering (CSB). This model introduces a trusted third party who is matching the needs of users with providers' offers. The involvement of CSBs is beneficial not only for customers, but also for providers. For instance, CSBs can attract more clients by suggesting the most appropriate services matching customers' needs. The presented problem lies in considering services which can be sold in bundles. Bundling is a common business practice, in which a set of services is sold together for the lower price than the sum of services' prices that are included in it.

This note introduces the grouping cloud services problem – a multi-criteria optimization problem which could help customers to determine the best choice of bundles (which contain cloud services) according to the chosen criteria.

4. Vincent Chau, Shenzhen Institutes of Advanced Technology, Chinese Academy of Sciences, China

Title: Competitive Algorithms for Demand Response Management in Smart Grid

Abstract: We consider a scheduling problem which abstracts a model of demand-response management in Smart Grid. In the problem, there is a set of unrelated machines and each job j (representing a client demand) is characterized by a release date, and a power request function representing its request demand at specific times. Each machine has an energy power function and the energy cost incurred at a time depends on the load of the machine at that time. The goal is to find a non-migration schedule that minimizes the total energy (over all times).

We give a competitive algorithm for the problem in the online setting where the competitive ratio depends (only) on the power functions of machines. In the setting with typical energy function $P(z) = z^\nu$, the algorithm is $\Theta(\nu^\nu)$ -competitive, which is optimal up to a constant factor. Our algorithm is robust in the sense that the guarantee holds for arbitrary request demands of clients. This enables flexibility on the choices of clients in shaping their demands — a desired property in Smart Grid.

We also consider a special case in offline setting in which jobs have unit processing time, constant power request and identical machines with energy function $P(z) = z^\nu$. We present a 2^ν -approximation algorithm for this case.

This is a joint work with Shengzhong Feng and Nguyen Kim Thang.

5. Xin Chen, Liaoning University of Technology, China

Title: Approximation Analysis on Flow-Shop Scheduling with Due Date Constraint

Abstract: In this talk, we consider the problem of scheduling in two-machine flow-shop systems to maximize total weighted early work, that is, the early part executed before jobs' common due date (F2|dj = d|max(Xw)). We mainly study the problem from the theoretical point of view. Firstly, we prove that any ordinary algorithm has an approximation ratio (at least) $2 \cdot \frac{w_{max}}{w_{min}}$, and Johnson algorithm has a lower bound of this value, indicating that it may be the “worst” one from the aspect of worst case analysis. Then, we give a FPTAS for the considered problem, which could be considered as the “best algorithms” from the theoretical point of view, and show the poor performances during its numerical experiments. We want to discuss that, is there any effective algorithm with good performance, both from theoretical and practical points of view.

Since we can hardly find the similar results under the considered subject (approximation analysis on flow-shop scheduling to optimize the criteria related with due date), we wonder the reasons for this phenomenon. (Maybe we miss them? Or the topic is not sufficient to be studied?)

6. Yong Chen, Hangzhou Dianzi University, China

Title: Scheduling Unit Jobs with Arbitrary Precedence Constraints on an m -Machine Flow Shop

Abstract: This paper studies flow shop scheduling with arbitrary precedence constraints that are defined by a directed acyclic graph (DAG). Each job (vertex) can start to be processed only when all its predecessor jobs in the DAG are finished. For unit-job scheduling on three machines to minimize makespan, a natural layered representation of the DAG leads to a $5/3$ -approximation algorithm. By improving the representation so that the most adjacent layers are connected by a pair of agreeable jobs, we propose a $3/2$ -approximation algorithm. We also claim that both algorithms can be applied to the general problem with m machines directly.

7. Yukun Cheng, Suzhou University of Science and Technology, China

Title: Agent Incentives of Strategic Behaviors in Resource Sharing on P2P Networks

Abstract: In this talk, we focus on the resource sharing over networks with autonomous participants (or agents), which goes beyond the peer-to-peer (P2P) bandwidth sharing idea. In such a resource sharing system, participants act as both suppliers and consumers of resources. Formally, this system is modeled as an undirected graph G , where each node i represents a peer with w_i units of divisible resources (or weight) to be distributed among its neighbors. Each agent obtains the utility by exchanging its resources with its neighbors according to the preset rules. However, agents may play strategically to improve their utilities by influencing the allocation, since the allocation depends on what they submit. Recently we proved the truthfulness of a market equilibrium mechanism, which can be obtained by a combinatorial method, against strategic behaviors of edge deletion and weight cheating. Another kind of strategy, called Sybil attack, is also considered. A strategic agent plays the Sybil attack by disguising itself to create several copied false nodes with its resources assigning among them. We are interested in the robustness of the market equilibrium mechanism in withstanding such a strategy in terms of the incentive ratio to measure how much one could gain at most. A series of work have been done on the Sybil attack strategy against the market equilibrium mechanism on different kinds of graphs.

8. Maciej Drozdowski, Poznan University of Technology, Poland

Title: Algorithm Portfolios for Large Berth Allocation Problems

Abstract: In this presentation algorithm selection for large berth allocation problem

(BAP) under solution time limits is considered. BAP consists in scheduling ships on berths subject to ready times, ship size constraints, for selected optimality criteria. In strategic port capacity planning large BAP instances must be solved many times in simulations and runtime allowance is an important simulation design decision. Sizes of instances which have to be considered in port capacity planning exceed by two orders of magnitude the sizes of BAP instances in the existing literature. Furthermore, BAP is **NP**-hard. Hence, selecting lightweight algorithms is essential.

Classic algorithm selection methods choose "one best" algorithm on basis of the central tendencies of algorithm performance on a set of test instances. We consider limitations of such methods. As an alternative an approach is proposed to select a portfolio of algorithms solving a BAP instance and returning the best solution. The algorithm portfolio must obey runtime limit while minimizing overall solution quality loss. Thus, the existing trade-off between algorithm runtime and solution quality is in the core of the portfolio selection method. The portfolio evolves with changing runtime limit which is a key design decision in the simulations. The portfolio is constructed on the basis of the performance of training instances measured by the runtime and solution quality.

Evolution of the algorithm portfolios under changing runtime limits is studied. Ability of the portfolios to solve new instances is evaluated. Performance of our portfolios is compared with other portfolios. For the training and validating instance datasets, random instances and real ship traffic logs are used.

9. Christoph Dürr, Sorbonne University, France

Title: Scheduling with a Processing Time Oracle

Abstract: Consider the problem of scheduling n given jobs on a single machine with the objective of minimizing the sum of the job completion times. Usually one needs to schedule the jobs in the order of increasing processing time. The novel difficulty in this problem is that the algorithm does not know the processing time of the jobs in beforehand. It only knows that they belong to the interval $[p, q]$. However the algorithm has the possibility to test a chosen job. The test is part of the schedule, takes one time unit, and reveals the actual job processing time to the algorithm. The performance of an algorithm is compared with the optimal schedule in terms of the so called competitive ratio. We show lower and upper bounds for this problem for two distinct models.

These are joint works with Fanny Dufossé, Thomas Erlebach, Julie Meißner, Nicole Megow, Noël Nadal, Denis Trystram and Óscar C. Vásquez.

10. Kan Fang, Tianjin University, China

Title: A Combinatorial Benders Decomposition for Parallel Machine Scheduling with Working-Time Restrictions

Abstract: This paper addresses an unrelated parallel machine scheduling problem with restrictions on employees' working-time and break time. In the problem, tasks must be processed nonpreemptively by employees on unrelated parallel machines, while different thresholds that specify the maximum total and consecutive working-time, and the minimum break time of each employee are respectively imposed. The objective is to minimize the weighted sum of the makespan, the machine depreciation costs and the labor costs. To solve this problem, a mixed integer linear programming model is formulated, and a combinatorial Benders decomposition algorithm and an LS-based heuristic are implemented. Extensive computational experiments are performed on randomly generated instances, and the results demonstrate the efficiency of our proposed combinatorial Benders decomposition approach.

11. Xin Han, Dalian University of Technology, China

Title: An Improved Speedup Factor for Sporadic Tasks with Constrained Deadlines under Dynamic Priority Scheduling

Abstract: Schedulability is a fundamental problem in real-time scheduling, but it has to be approximated due to the intrinsic computational hardness. As the most popular algorithm for deciding schedulability on multiprocess platforms, the speedup factor of partitioned-EDF is challenging to analyze and is far from been determined. Partitioned-EDF was first proposed in 2005 by Barush and Fisher, and was shown to have a speedup factor at most $3-1/m$, meaning that if the input of sporadic tasks is feasible on m processors with speed one, partitioned-EDF will always return succeeded on m processors with speed $3-1/m$. In 2011, this upper bound was improved to $2.6322-1/m$ by Chen and Chakraborty, and no more improvements have appeared ever since then. In this paper, we develop a novel method to discretize and regularize sporadic tasks, which enables us to improve, in the case of constrained deadlines, the speedup factor of partitioned-EDF to $2.5556-1/m$, very close to the asymptotic lower bound 2.5.

12. Tamás Kis, Acamdey of Sciences, Hungary

Title: Machine Scheduling with non-Renewable Resources

Abstract: None-renewable resources, like raw materials, energy, on money, share the common property that their availability is decreased by consuming them. This feature

is in strong contrast with renewable resources, like machines, CPUs, human operators, tools, which are only reserved by the jobs during their processing. In the talk, I overview machine scheduling problems, where in addition to a single, or parallel machines, there are additional non-renewable resources consumed by the jobs. Each non-renewable resource has an initial stock, and some replenishments over time at given dates and in known quantities. I consider only off-line, deterministic problems, and the main question is the approximability of problems in this class with various objective functions. In particular, the approximability of the makespan objective in the single machine, and in the identical parallel machine environment is almost fully understood, and I show the most important tricks in the design of approximation schemes (mostly PTAS, or sometimes FPTAS) for these problems. Much less is known about the maximum lateness objective, where additional assumptions are needed to reach matching results. Finally, I overview the recent finding in the approximability of the total weighted completion time objective in a single machine environment. As we will see, the latter objective function is far more challenging to approximate than the makespan.

13. Jing Li, New Jersey Institute of Technology, USA

Title: New Parallel Scheduling Problems in the Presence of GPU Accelerators

Abstract: Many emerging systems today, including self-driving cars, unmanned aerial vehicles, robotics, and various intelligent cloud services, heavily rely upon artificial intelligence (AI) and machine learning (ML) components to perform critical operations. Since these AI/ ML components have high data parallelism and are computationally intensive, graphics processing units (GPUs) are coming to be widely adopted in systems as computing accelerators for such components. This trend makes it imperative that we better understand how to schedule workloads upon CPU+GPUs. This talk will provide an abstraction of the workloads and computational platforms that contain a GPU as an accelerator, under which challenging new scheduling problems arise. We will further discuss the connection between the new problems and the traditional scheduling problems, such as flow shop and parallel job scheduling, and present our preliminary results exploiting such connection.

14. Minming Li, The City University of Hong Kong, China

Title: Truthful Mechanisms for Location Games of Dual-Role Facilities

Abstract: This paper is devoted to the facility location games with payments, where every agent plays a dual role of facility and customer. In this game, each selfish agent is located on a publicly known location in a metric space, and can allow a facility to

be opened at his place. But the opening cost is his private information and he may strategically report this opening cost. Besides, each agent also bears a service cost equal to the distance to his nearest open facility. We are concerned with designing truthful mechanisms for the game, which, given agents' reports, output a set of agents whose facilities could be opened, and a payment to each of these agents who opens a facility. The objective is to minimize (exactly or approximately) the social cost (the total opening and service costs) or the maximum agent cost of the outcome.

We characterize the normalized truthful mechanisms for this game. Concerning the minimum social-cost objective, we give an optimal truthful mechanism without regard to time complexity, and show a small gap between the best known approximation ratio of polynomial-time truthful mechanisms for the game and that of polynomial-time approximation algorithms for the counterpart of pure optimization. For the minimum maximum-cost objective, we provide an optimal truthful mechanism which runs in polynomial time. We also investigate mechanism design for the game under a budget on the total payment.

15. Rongheng Li, Hunan Normal University, China

Title: Nash Equilibrium Analysis for Scheduling Games on Identical Machines

Abstract: Given a set n tasks and a set of m identical machines, in which tasks and machines are possessed by different selfish clients. Each selfish client of machine gets a profit equal to its load and each selfish client of task suffers from a cost equal to the load of the machine on which it is allocated. Our aim is to allocate the tasks on the m machines so as to minimize the maximum completion times of the tasks on each machine. A stable allocation is referred to as a dual equilibrium (DE). We firstly show that $4/3$ is the tight upper bound of the price of anarchy(PoA) with respect to dual equilibrium for $3 \leq m \leq 9$. Secondly we address that $(7m-6)/(5m-3)$ is an upper bound for $m \geq 10$. The result is better than the existing bound of $7/5$.

16. Wenhua Li, Zhengzhou University, China

Title: Approximation Algorithms for Parallel-Batch Scheduling with Processing Set Restrictions

Abstract: We consider job scheduling on m parallel-batch machines to minimize the makespan. Each job has a given set of machines to be assigned. Each machine can process several jobs simultaneously as a batch. We study two models: (i) scheduling on parallel-batch machines with the nested processing set restrictions and (ii) scheduling on uniform parallel-batch machines with the tree-hierarchical processing set restrictions.

The uniform parallel-batch machines have different speeds. For the first model, we design a strongly polynomial-time solution algorithm with a worstcase bound of 2 for the case with non-identical batch capacity and a worst-case bound of $2-1/m$ for the case with identical batch capacity. It is better than a known polynomial-time algorithm with a worst case bound of $3-1/m$, which only applies to the case with identical batch capacity. For the second model, we design a fast algorithm with a worst-case bound of 2 for the case with non-identical batch capacity. It is better than a known fast algorithm with a worst-case bound of $9/4$, which only applies to the case with parallel-batch machines and the inclusive processing set restrictions. We also provide a fast algorithm with a worst-case bound of $2-1/m$ for the case with identical batch capacity.

17. Yuanzheng Li, Huazhong University of Science and Technology, China

Title: Flexible Scheduling of Microgrid with Uncertainties Considering Expectation and Robustness

Abstract: With increasing renewable distributed energy resources integrated into a microgrid (MG), it is difficult to realize its optimal operation due to various of uncertainties, e.g., power outputs of wind turbines and photovoltaic cells. This paper presents a flexible scheduling model for the optimal operation of a MG, in which the expectation of operational cost and the robustness of scheduling solutions are considered, simultaneously. To balance the trade-off between minimizing the expectation of operational cost and enhancing robustness, the proposed flexible scheduling model of MG is formulated as a multi-objective optimization problem, by introducing a robustness preference parameter. In this way, the relationship between the expected operational cost and the robustness of scheduling solution is investigated, and the fuzzy decision making method is used for selecting a suitable solution. Finally, a modified MG is used for case study, and simulation results verify the applicability and effectiveness of the proposed model.

18. Lingfa Lu, Zhengzhou University, China

Title: In-house Production Scheduling and Outsourcing under Different Discount Schemes

Abstract: In this paper, we consider a coordinated in-house production scheduling and outsourcing model, where a manufacturer might outsource part of the production to a subcontractor. The manufacturer pays a specific outsourcing cost for each job that is outsourced. To encourage the manufacturer to outsource more jobs, the subcontractor provides a specific discount scheme, which depends on the total outsourcing cost or on the number of outsourced jobs. Our model is the first to investigate the impact on the

manufacturer's decision-making under different discount schemes. Some usual distinct discount schemes are studied. For each problem, we either provide a polynomial-time exact algorithm, or show that the problem is NP-hard and then propose effective approximation algorithms.

19. Wenchang Luo, Ningbo University, China

Title: An Approximation Scheme for Rejection-Allowed Rescheduling on a Single Machine with Job Delays

Abstract: We study a novel rescheduling problem in which a set of jobs has been assigned an original schedule to minimize the total weighted completion time on a single machine, with the assumption that all of them will be available at the planned beginning of the processing. The need for rescheduling arises due to some jobs could not arrive in time; the decision-maker has to adjust the original schedule to account for the delayed jobs without causing excessive time disruption to the original schedule and to minimize his/her operational cost. While the decision-maker can choose to reject some jobs, the tardiness of each accepted job in the adjusted schedule is strictly upper bounded, for which the completion time of the job in the original schedule is taken as its due date. The total operational cost includes three components: the total weighted completion time of the accepted jobs, the total rejection cost of the rejected jobs, and the penalty on the maximum tardiness for the accepted jobs. We study this novel rescheduling problem from approximation algorithm perspective, as it generalizes several classic NP-hard scheduling problems, design a pseudo-polynomial time dynamic programming exact algorithm, and furthermore, using the standard sparsifying technique, develop the exact algorithm into a fully polynomial time approximation scheme.

20. Benjamin Moseley, Carnegie Mellon University, USA

Title: Online Scheduling via Learned Weights

Abstract: The use of machine learning and predictive models have produced a revolution in science and engineering. Online optimization problems are a natural source of uncertainty that predictions can be used to manage and improve performance. This paper studies how predictive techniques can be used to break through worst case barriers in online scheduling.

The make-span minimization problem on unrelated machines and its special case, restricted assignment, are classic problems in online scheduling theory. Worst case analysis of these problems yields $\Omega(\log m)$ lower bounds on the competitive ratio in the online setting. In this paper we construct non-trivial predictions for these problems

and design algorithms that utilize these predictions to compute solutions online. Our predictions are compact in size, having dimension linear in the number of machines. The performance guarantees of our algorithms depend on the accuracy of the predictions, and moderately accurate predictions allow our techniques to beat the worst case lower bounds. More precisely, the predictions can be used to construct $O(\log \eta)$ -competitive fractional assignments, where η is the error of the predictions. We then give an online algorithm that is $O((\log \log(m))^3)$ -competitive to round these fractional assignments, with a nearly matching $\Omega(\log \log m)$ lower bound for online rounding algorithms.

21. Kameng Nip, Sun Yat-sen University, China

Title: Two-Machine Flow Shop Scheduling Problem under Linear Constraints

Abstract: In this talk, we introduce a two-machine flow shop scheduling problem under linear constraints (2-FLC problem in short), in which the processing times of two stages of jobs are also decision variables and satisfy a system of linear constraints. The goal is to determine the processing times of each job, and to schedule the jobs to the two-machine flow shop such that the makespan, i.e., the completion time of all the jobs is minimized. This problem can find application in various areas, such as industrial production and planning. We study the computational complexity and algorithms for the 2-FLC problem. Particularly, we show that although the two-machine flow shop scheduling problem can be solved in polynomial time, the 2-FLC problem is generally NP-hard in the strong sense. Then we consider the design of algorithms on various setting of the 2-FLC problem. In particular, we design a polynomial time algorithm for the 2-FLC problem when there is a fixed number of constraints. For the general case, we first propose a simple 2-approximation algorithm, and then design a polynomial time approximation scheme (PTAS). Finally, we extend the result to the 2-FLC problem with no-wait constraint and no missing operations. We show that this case is also NP-hard, and design algorithms for various settings.

22. Qingqin Nong, Ocean University of China, China

Title: Maximize a Monotone Function with a Generic Submodularity Ratio

Abstract: Generic submodularity ratio is a general measurement to characterize how close a nonnegative monotone set function is to be submodular. In this paper we make a systematic analysis of greedy algorithms for maximizing a monotone and normalized set function with a generic submodularity ratio under Cardinality constraints, Knapsack constraints, Matroid constraints and K-intersection constraints.

23. Franciszek Seredynski, University UKSW, Warsaw, Poland

Title: Load Balancing and Scheduling Problems: an Approach Based on Spatial Prisoner's Dilemma Game

Abstract: First we study conditions of emergence of collective behavior of agents acting in the 2D space, where each agent takes part in Prisoner's Dilemma (PD) game with opponents from a neighborhood. We use two classes of agents: cellular automata and learning automata. While each agent is oriented on maximization of its own profit, we are interested in answering the question if and when a phenomenon of global cooperation in a large set of agents is possible. We present results of the experimental study showing conditions and degree of emerging such a cooperation. Next, we apply this approach to solve the problem of security-aware scheduling and load balancing in Cloud Computing systems. Brokers representing users participate in the game to fulfill their own three criteria: the execution time of the submitted tasks, their execution cost and the level of provided security assurance. We experimentally show that in the process of the game a solution is found which provides an optimal resource utilization while users meet their applications' performance and security requirements with minimum expenditure and overhead. In the last part of the talk we apply an approach based on the spatial PD game to solve the Maximal Lifetime Coverage Problem in Wireless Sensor Networks.

24. Jiri Sgall, Charles University, Czech

Title: A ϕ -Competitive Algorithm for Scheduling Packets with Deadlines

Abstract: In the *online packet scheduling problem with deadlines*, the goal is to schedule transmissions of packets that arrive over time in a network switch and need to be sent across a link. Each packet p has a deadline d_p , representing its urgency, and a non-negative weight w_p , that represents its priority. Time is discrete, with time units represented by consecutive integers that we refer to as *time slots* or *steps*. Only one packet can be transmitted in any time slot. The objective is to compute a schedule that maximizes the total weight of packets which are successfully transmitted. In the literature this problem is also occasionally referred to as *bounded-delay buffer management*, *QoS buffering*, or as a job scheduling problem for unit-length jobs with release times, deadlines, and weights, where the objective is to maximize the weighted throughput.

The problem is inherently online, with the scheduling decisions made without the knowledge of future packet arrivals. The central problem, intensively studied since 2001, is to determine the optimal competitive ratio of online algorithms, namely the worst-case ratio between the optimum total weight of a schedule (computed by an

offline algorithm) and the weight of a schedule computed by a (deterministic) online algorithm. We solve this open problem by presenting a ϕ -competitive online algorithm (where $\phi \approx 1.618$ is the golden ratio), matching the previously established lower bound.

This result is a joint work with Pavel Veselý, Marek Chrobak, and Łukasz Jeż. It appears in the proceedings of SODA 2019 and in arXiv:1807.07177.

25. Gaurav Singh, BHP Billiton, Australia

Title: Multi-Criteria Optimisation of Ramp Designs in Open-Pit Mining

Abstract: The optimal mine plans developed by sophisticated algorithms are transformed into an operational plans by developing ramps. These ramps are used by mining equipments like trucks and shovels to transport the ore to destinations on the surface. The optimal ramp design requires to trade-off three different criteria, operating cost of moving the material, capital cost of building these ramps and loss of revenue due to sterilisation of mining area below the ramps. This talk will present a methodology for modelling such problems and also a multi-criteria algorithm for solving these large scale problems. Results from real mining operations will also be presented.

26. Abhinav Srivastav, University of Evry, France

Title: Online non-Preemptive Scheduling to Minimize Maximum Weighted Flow-Time on Related Machines

Abstract: We consider the problem of scheduling jobs to minimize the maximum weighted flow-time on a set of related machines. When jobs can be preempted this problem is well-understood; for example, there exists a constant competitive algorithm using speed augmentation. When jobs must be scheduled non-preemptively, only hardness results are known. In this paper, we present the first online guarantees for the non-preemptive variant. We present the first constant competitive algorithm for minimizing the maximum weighted flow-time on related machines by relaxing the problem and assuming that the online algorithm can reject a small fraction of the total weight of jobs. This is essentially the best result possible given the strong lower bounds on the non-preemptive problem without rejection.

27. Rob van Stee, University of Siegen, Germany

Title: The Price of Clustering in Bin-Packing with Applications to Bin-Packing with Delays

Abstract: One of the most significant algorithmic challenges in the “big data era” is

handling instances that are too large to be processed by a single machine. The common practice in this regard is to partition the massive problem instance into smaller ones and process each one of them separately. In some cases, the solutions for the smaller instances are later on assembled into a solution for the whole instance, but in many cases this last stage cannot be pursued (e.g., because it is too costly, because of locality issues, or due to privacy considerations). Motivated by this phenomenon, we consider the following natural combinatorial question: Given a *bin-packing* instance (namely, a set of items with sizes in $(0, 1]$ that should be packed into unit capacity bins) I and a partition $\{I_i\}_i$ of I into *clusters*, how large is the ratio $\sum_i \mathbf{Opt}(I_i)/\mathbf{Opt}(I)$, where $\mathbf{Opt}(J)$ denotes the optimal number of bins into which the items in J can be packed?

In this paper, we investigate the supremum of this ratio over all instances I and partitions $\{I_i\}_i$, referred to as the bin-packing *price of clustering* (PoC). It is trivial to observe that if each cluster contains only one tiny item (and hence, $\mathbf{Opt}(I_i) = 1$), then the PoC is unbounded. On the other hand, a relatively straightforward argument shows that under the constraint that $\mathbf{Opt}(I_i) \geq 2$, the PoC is 2. Our main challenge was to determine whether the PoC drops below 2 when $\mathbf{Opt}(I_i) > 2$. In addition, one may hope that $\lim_{k \rightarrow \infty} \text{PoC}(k) = 1$, where $\text{PoC}(k)$ denotes the PoC under the restriction to clusters I_i with $\mathbf{Opt}(I_i) \geq k$. We resolve the former question affirmatively and the latter one negatively: Our main results are that $\text{PoC}(k) \leq 1.951$ for any $k \geq 3$ and $\lim_{k \rightarrow \infty} \text{PoC}(k) = 1.691\dots$. Moreover, the former bound cannot be significantly improved as $\text{PoC}(3) > 1.933$. In addition to the immediate contribution of this combinatorial result to “big data” kind of applications, it turns out that it is useful also for an interesting online problem called *bin-packing with delays*.

28. Malgorzata Sterna, Poznan University of Technology, Poland

Title: Parallel Machine Scheduling Problems with Late Work and Early Work

Abstract: Scheduling problems with the criteria related to late work have been investigated for more than 30 years [1]. Late work is an objective function based on parts of jobs executed after their due dates, while the complementary performance measure - early work - is based on parts of jobs executed before their due dates. Both functions find many practical applications, for example in control systems, manufacturing systems or agriculture. Late work based criteria have received much more attention than early work based ones so far. At the beginning, early work was considered as an auxiliary parameter, simplifying notation in the analysis of late work related functions, and it has become the separate subject of research quite recently [2, 8]. Both criteria are related to two other scheduling domains: scheduling imprecise computations and

scheduling with controllable processing times.

Within the talk, the results of studies on the basic scheduling model with parallel identical machines and a common due date are collected for criteria based on late work, as well as on early work. They include the problem complexity analysis and proposals of exact and approximation algorithms, as well as online algorithms, and proposals of conventional search approaches. Moreover some more general studies on the interrelations among late and early work based criteria are mentioned [4, 9].

First, the binary NP-hardness for two machines and unary NP-hardness for at least three machines was proven [6]. For the two-machine model, the transformation from the partition problem and the pseudo-polynomial time dynamic programming were given. These results were extended for the model with the fixed number of machines [5]. For the problem with the arbitrary number of machines, the transformation from the three-partition problem was showed. The complexity proofs and the proposals of exact algorithms are valid both for the late work minimization, as well as for early work maximization. However due to the specificity of late work, which can be minimized to zero for some problem instances, the approximation algorithms can be proposed only for the problems with early work maximization [9]. First, the approximation algorithm based on the list strategy, adopted from the bin packing problem, was proposed with the approximation ratio depending on the number of machines [6]. Moreover, the worst case performance of the classical list method based on the longest processing time rule, was estimated [7]. Then, for the two machine problem, the polynomial time approximation scheme was proposed [9]. Recently, the fully polynomial time approximation scheme for the more general problem with the fixed number of machines was formulated [5], which is obviously also valid for the two-machine case.

The heuristic algorithm mentioned before, adopted from the bin packing problem, was also applied to solve an online version of the parallel machine problem, where jobs appear in the system one by one, and a current job has to be scheduled before a new one comes without the knowledge on the whole problem instance (so called online over list model). The competitive analysis and the proof of the problem lower bound allowed determining the performance of this approach in online mode [6]. Then, the semi-online versions of the problem were studied, where partial information on the problem instance, such as the maximum job processing time or/and the optimal value of the criterion function, is given in advance at the beginning of the scheduling process. A few semi-online algorithms were proposed for these models with the constant competitive ratios [3].

The theoretical studies mentioned above were completed with some computational

experiments performed for list scheduling heuristics for the two machine problem [9], as well as for dynamic programming and approximation schemes for the problem with the fixed number of machines [5]. Moreover, for the more general model with the arbitrary job due dates, the branch and bound method and the genetic algorithm were proposed [7], which can be obviously applied for the common due date case too.

References

- [1] Blazewicz J. Scheduling preemptible tasks on parallel processors with information loss. *Technique et Science Informatiques* 3(6), 415–420, 1984.
- [2] Blazewicz J, Ecker K, Pesch E, Sterna M, Schmidt G, Weglarz J. Handbook on scheduling. From theory to practice. 2nd ed. Chapter 14: Scheduling Imprecise Computations, 527–576, Berlin-Heidelberg-New York: Springer; 2019.
- [3] Chen X, Kovalev S, Liu Y, Sterna M, Chalamon I, Blazewicz J. Semi-online scheduling on two identical machines with a common due date to maximize total early work. *Technical Report RA-3/2019*, Institute of Computing Science, Poznan University of Technology, 2019.
- [4] Chen X, Kovalev S, Sterna M, Blazewicz J. Mirror scheduling problems with early work and late work criteria. *Technical Report RA-1/2018*, Institute of Computing Science, Poznan University of Technology, 2019.
- [5] Chen X, Liang Y, Sterna M, Wang W, Blazewicz J. Fully polynomial time approximation scheme to maximize early work on parallel machines. *Technical Report RA-2/2019*, Institute of Computing Science, Poznan University of Technology, 2019.
- [6] Chen X, Sterna M, Han X, Blazewicz J. Scheduling on parallel identical machines with late work criterion: offline and online cases. *Journal of Scheduling* 19(6), 729–736, 2016.
- [7] Chen X, Wang W, Xie P, Zhang X, Sterna M, Blazewicz J. Exact and approximation algorithms for scheduling on two identical machines with early work maximization. *Technical Report RA-4/2019*, Institute of Computing Science, Poznan University of Technology, 2019.
- [8] Sterna M. A survey of scheduling problems with late work criteria. *Omega* 39(2), 120–129, 2011.

- [9] Sterna M, Czerniachowska K. Polynomial time approximation scheme for two parallel machines scheduling with a common due date to maximize early work. *Journal of Optimization Theory and Applications* 174(3), 927–944, 2017.

29. Vitaly Strusevich, Greenwich University, UK

Title: Parallel Jobs of Two Sizes: Preemptive Scheduling of with Controllable Processing Times

Abstract: In parallel machine scheduling, a size of a job is defined as the number of machines that are simultaneously required for its processing. This paper considers a scheduling problem in which the jobs have two sizes. The processing times are controllable and they have to be chosen for given intervals in order to guarantee the existence of a preemptive schedule in which all jobs are completed by a common deadline. The objective is the total compression cost which reflects possible reductions in processing times. Unlike problems of classical scheduling, with jobs of size one, the model under consideration cannot be directly handled by submodular optimization methods. We develop a fast polynomial-time algorithm.

Formally, in scheduling on (identical) parallel machines, we are given a set $N = \{1, 2, \dots, n\}$ of jobs/tasks to be processed on $m \geq 2$ machines/processors M_1, M_2, \dots, M_m . It takes $p(j)$ time units to process a job $j \in N$. Preemption is allowed, i.e., the processing of any job j can be interrupted and resumed later; the total length of time intervals of processing job j must be equal to $p(j)$. In the classical setting, each job can be processed preemptively on any machine, but only one at a time. In the model with rigid jobs, for each job $j \in N$ a number of machines, denoted by $size(j)$, where $1 \leq size(j) \leq m$, is given and at any time of its processing the job must be preemptively processed on exactly $size(j)$ machines simultaneously. We call jobs of size larger than 1 *parallel jobs* or *multi-processor jobs*. See the monograph [2] for a comprehensive exposition of scheduling with multi-processor jobs.

The main scheduling model studied in this paper, is a restricted model with rigid jobs, in which conventional jobs are processed by one machine and all parallel jobs have the same size Δ , where $1 < \Delta \leq m$.

For a feasible schedule, let the completion time of job $j \in N$ be denoted by $C(j)$. Assuming that the processing times $p(j)$ are *fixed*, the following *feasibility* problem is of interest: to check whether there exists a schedule in which all jobs are completed by a given deadline D . For the model with parallel jobs of two sizes, the feasibility problem can be solved in $O(n)$ time; see [1]. Notice that if the sizes of jobs are arbitrary, the

problem with parallel jobs is still solvable in $O(n)$ time, provided that the number of machines m is fixed; see [3].

In the models of scheduling with controllable processing times (the SCPT models) the actual processing time $p(j)$ of job $j \in N$ is not given in advance but has to be chosen by a decision-maker from a given interval $[l(j), u(j)]$. Such a decision results in *compression* of the longest processing time $u(j)$ down to $p(j)$, and the value $x(j) = u(j) - p(j)$ is called the *compression amount* of job j . Compression may decrease the completion time of each job j but incurs additional cost, denoted by $w(j)$.

A wide range of the SCPT problems has been studied. For problems relevant to this paper, the goal is to minimize the total compression cost $\sum_{j \in N} w(j)x(j)$ for a schedule in which all jobs complete by a common deadline D . See the survey [4] and the references within for the best known polynomial-time algorithms for these and more general problems, that allow job-dependent release dates and/or deadlines and machine speeds. All previously known results on the SCPT problems, however, handle the models of the classical type, i.e., those with the jobs of size one. These techniques cannot be directly extended to handle with problems with parallel jobs, and an alternative method is used.

The set N of jobs is split into two subsets: N_Δ of Δ -jobs, i.e., the jobs of size Δ , and N_1 of 1-jobs, i.e., the jobs of size 1. Let $p(1), p(2), \dots, p(n)$ be the variables that represent the values of the actual processing times such that $l(j) \leq p(j) \leq u(j)$ for each $j \in N$.

We introduce a parameter λ , equal to total work of the Δ -jobs, i.e.,

$$\lambda = \sum_{j=1}^{|N_\Delta|} p(j).$$

Provided that total work of Δ -jobs is equal to λ and the deadline D is observed, introduce the following functions of parameter λ :

- $W_\Delta(\lambda)$, the maximum total weighted work of Δ -jobs;
- $W_1(\lambda)$, the maximum total weighted work of 1-jobs;
- $W(\lambda)$, the maximum total weighted work of all jobs.

The structure of the problem allows us to employ the following approach to finding $W(\lambda)$. First, we consider the Δ -jobs alone and find their maximum total work $W_\Delta(\lambda)$, as a piecewise-linear concave function of λ . Then, keeping the structure of the corresponding schedule for processing the Δ -jobs, we add the 1-jobs and maximize $W_1(\lambda)$.

The running time of the resulting algorithm that determines $W(\lambda) = W_{\Delta}(\lambda) + W_1(\lambda)$ is $O(n \log n)$ time.

References

- [1] Błażewicz, J., Drabowski, M., Węglarz, J. : Scheduling multiprocessor tasks to minimize schedule length. IEEE Trans. Comput. C-35, 389–393 (1986)
- [2] Drozdowski, M.: Scheduling for Parallel Processing. Springer, London (2009)
- [3] Jansen, K., Porkolab, L.: Computing optimal preemptive schedules for parallel tasks: linear programming approaches. Math. Progr. A, 95, 617–630 (2003)
- [4] Shioura, A., Shakhlevich, N.V., Strusevich V.A.: Preemptive models of scheduling with controllable processing times and of scheduling imprecise computation: A review of solution approaches. Eur. J. Oper. Res. 266, 795-818 (2018)

30. Nguyen Kim Thang, University Paris Saclay, France

Title: An Improved Approximation Algorithm for Scheduling under Arborecence Precedence Constraints

Abstract: We consider a scheduling problem on unrelated machines with precedence constraints. There are m unrelated machines and n jobs and every job has to be processed non-preemptively in some machine. Moreover, jobs have precedence constraints; specifically, a precedence constraint $j < j'$ requires that job j' can only be started whenever job j has been completed. The objective is to minimize the total completion time.

The problem has been widely studied in more restricted machine environments such as identical or related machines. However, for unrelated machines, much less is known. In the paper, we study the problem where the precedence constraints form a collection of arborecences. We present a $O((\log n)^2/(\log \log n)^3)$ -approximation algorithm — that improves the best-known guarantee of $O((\log n)^2/\log \log n)$ due to Kumar et al.’ 09 a decade ago. The analysis relies on a dual-fitting method in analyzing the Lagrangian function of non-convex programs.

31. Denis Trystram, Grenoble Institute of Technology, France

Title: New Challenges of Scheduling at the Edge

Abstract: Edge Computing is a paradigm that emerged recently. The concept is to compute as close as possible to where the data are produced in order to avoid too heavy data movements. Typical scheduling algorithms have been designed for classical computing platforms are not suitable for the edge. We will first discuss why (the main characteristics of edge include high heterogeneity and volatility of the computing components, high latency and uncertainty of the network connections). We will describe a scheduling problem and present a scheduling algorithm. This work is based on a collaboration with the company Qarnot Computing company, which operate a private cloud composed of smart heaters (built upon multi-cores). We will finally report some experimental results obtained on a simulator.

This work has been done in common with Clement Mommessin, Yanik Ngoko and Giorgio Lucarelli.

32. Marc Uetz, University of Twente, The Netherlands

Title: News on Unrelated Machine Scheduling under Uncertainty

Abstract: This talk reports about an update of a 2017 IPCO paper together with Varun Gupta, Ben Moseley, and Qiaomin Xie. In that updated paper, we show that a rather simple and intuitive greedy algorithm performs surprisingly well for a classical scheduling problem, namely minimizing the total weighted completion time on unrelated machines. In fact we give the first results for this problem when jobs are allowed to appear online over time, and have uncertain job sizes. Due to recent simplifications and improvements in both algorithm and analysis, our new performance bounds even improve upon previously best known results for the deterministic online problem, from the 1990s. The algorithm’s basic idea is a greedy assignment of jobs to machines, just mimicking a “nominal” schedule where stochastic processing times are replaced by expectations. Moreover, the main idea for making this algorithm also competitive for uncertain job sizes, is to adhere as much as possible to that nominal schedule. The analysis is based on dual fitting.

33. Yinling Wang, Dalian University of Technology, China

Title: A New Approximation Algorithm for Flow Shop with Transporter Coordinate

Abstract: In this paper, we study a problem of the two-machine flow shop scheduling problem with intermediate transportation. The best approximation algorithm was presented with a two approximation ratio to our best knowledge. In this paper, We propose a $(5/3 + \epsilon)$ -approximation algorithm for this problem, where $\epsilon > 0$. Moreover, our algorithm can reach the lower bound $5/3$ asymptotically given when ϵ is close to 0.

34. Deshi Ye, Zhejiang University, China

Title: Quadratic Knapsack Problems with Scheduling Applications

Abstract: Both the knapsack problem and the scheduling problem are traditional combinatorial optimization problems. In this talk, we first introduce how various scheduling problems can be modeled as quadratic knapsack problems dealing with weight and value constraints. In the second part, we design approximation algorithms to solve a multiple quadratic knapsack problem. When the given profit matrix is positive semidefinite and of constant rank, we present an FPTAS for the single knapsack problem, and a PTAS for multiple m knapsacks, where the number of knapsacks m is a constant.

35. Wei Yu, East China University of Science and Technology, China

Title: Exact and Approximation Algorithms for Some Variants of the Min-Max k -Traveling Salesman Problem on a Tree

Abstract: Given an edge-weighted tree $T=(V,E)$ with a depot $s \in V$ and k identical salesmen, where k is independent of the input, the min-max k -traveling salesmen problem on a tree (Min-Max k -TSPT) is to find a set of k tours, each of which starting from and ending at s , for the salesmen to serve all customers located at the vertices of the tree such that the maximum total edge weight of the tours is minimized. In this paper, we present new simplified pseudo-polynomial time exact algorithms for both the Min-Max k -TSPT and its multi-depot extension, where there is a candidate depot set $D \subseteq V$ and each tour is required to start from and end at the same depot in D . Based on this simplified approach, we devise pseudo-polynomial algorithms for some natural variants of the Min-Max k -TSPT and its multi-depot generalization. Moreover, we transform these pseudo-polynomial algorithms into FPTASes by standard rounding technique.

36. Xiandong Zhang, Fudan University, China

Title: Scheduling of Coupled Tasks with Exact Delays for Minimum Total Job Completion Time

Abstract: In this paper, we fill in a conspicuous gap in research on scheduling coupled tasks. We draw a full complexity picture for single-machine scheduling of coupled tasks of exact time delays in between with the objective of minimizing the total of job completion times. Three basic problems are shown to be NP-hard in the strong sense. Polynomial time algorithms are provided for those tractable cases. This basic model of scheduling coupled tasks has been studied for various radar racking systems, and also been used for scheduling systems of robotic cells and patient appointments in a

chemotherapy outpatient clinic.

37. Yong Zhang, Shenzhen Institutes of Advanced Technology, Chinese Academy of Sciences, China

Title: Approximation Algorithm and Incentive Ratio of the Selling with Preference

Abstract: We consider the market mechanism to sell two types of products, A and B, to a set of buyers $I = 1, 2, \dots, n$. The amounts of products are m_A and m_B respectively. Each buyer i has his information including the budget, the preference and the utility function. On collecting the information from all buyers, the market maker determines the price of each product and allocates some amount of product to each buyer. The objective of the market maker is design a mechanism to achieve the semi market equilibrium. In this paper, we show that maximizing the total utility of the buyers in satisfying the semi market equilibrium is NP-hard and give a 1.5-approximation algorithm for this optimization problem. Moreover, in the market, a buyer may get more utility by misreporting his information. We consider the situation that a buyer may misreport his preference and prove that the incentive ratio, the percentage of the improvement by misreporting the information, is upper bounded by 1.618.