04231 Abstracts Collection
Scheduling in Computer and Manufacturing Systems
— Dagstuhl Seminar —

J. Błażewicz[1], K. Ecker[2], E. Pesch[3] and D. Trystram[4]

[1] Politechnka Poznanska, Instytut Informatyki
ul. Piotrowo 3A, 60-965 Poznan, PL
jblazewicz@cs.put.poznan.pl
[2] TU Clausthal, Institut für Informatik
Julius-Albert-Str. 4, 38678 Clausthal-Zellerfeld, D
ecker@informatik.tu-clausthal.de
[3] Universität Siegen
FB 5 - Wirtschaftswissenschaften, LST Wirtschaftsinformatik
Hölderlinstr. 3, 57068 Siegen, D
pesch@fb5.uni-siegen.de
[4] ID-IMAG
51 rue Jean Kuntzman, 38330 Montbonnot St. Martin, FR
denis.trystram@imag.fr

Abstract. During 31.05.-04.06.04, the Dagstuhl Seminar 04231 "Scheduling in Computer and Manufacturing Systems" was held in the International Conference and Research Center (IBFI), Schloss Dagstuhl. During the seminar, several participants presented their current research, and ongoing work and open problems were discussed. Abstracts of the presentations given during the seminar as well as abstracts of seminar results and ideas are put together in this paper. The first section describes the seminar topics and goals in general. Links to extended abstracts or full papers are provided, if available.

Keywords. Scheduling

Overview

The biannual Workshop on "Scheduling in Computer and Manufacturing Systems" was organized in Dagstuhl Castle during the week May 31-June 4, 2004. It was the fifth meeting and gathered 58 participants from universities or research centers in Austria, Belarus, Belgium, Brazil, Canada, France, Germany, Great Britain, Hong Kong, Hungary, Israel, Italy, Poland, Switzerland, The Netherlands, and the U.S.A.

The objective of the seminar was to provide a forum for the discussion of ongoing research in scheduling. The seminar promoted an exchange of ideas covering the entire spectrum from case studies of real applications to recent advances

in mathematical foundations. The various aspects of scheduling were covered by 39 lectures that addressed classical application areas such as distributed processing, operating systems, dependable systems, and flexible manufacturing. It is worth pointing out that many lectures were motivated by practical considerations, as for example machine break-downs, batch scheduling, synchronous production, robotic cell scheduling, real-time scheduling, and resource investment problems. But also exciting new areas emerged such as those in modern communications systems, examples being wireless networks, multimedia networks, and the internet.

The seminar proceeded along three broad fronts: applications, which include empirical studies of existing systems as well as numerical studies of the analysis and simulation of system models. Most of the application studies came from the area of production scheduling and planning, such as just-in-time scheduling, due date assignment and project control, including special problems dealing with machine break-downs, robotic cells, assembly scheduling, load balancing, minimizing the number of workers (human resources). Other presentations considered special problems from chemistry and oceanography, the design of schedulers, e.g. for web applications, and planning examination sessions. Algorithms were presented for various problems such as batch scheduling, resource scheduling, tardiness problems, shop problems, deadline and due date scheduling, real-time scheduling, on-line scheduling, single machine problems, time lags, scheduling with communication delays. The main concern in these presentations was the design and analysis of algorithms ranging from simple and tractable on-line and greedy rules to methods based on semi-enumerative approaches, branch and bound, local neighborhood search, and LP formulations. New theoretical developments included recent results in the analysis of new and classical problems under novel (or multiple) criteria, dealing with particular assumptions on machines, tasks (e.g., release dates, precedence constraints, communication delays, multiprocessor tasks, bi-processor tasks), and other problems such as assembly scheduling problems and on-line scheduling. Typical questions discussed were the structure of problems and their relation to graph theory, complexity of problems including polynomial solvability, the design of algorithms and performance analysis, and the approximability of optimal solutions.

The participants were delighted with the outstanding local organization and the marvelous facilities that created the atmosphere for a successful seminar.

On behalf of the participants,
J. Błażewicz, E. Pesch, K. Ecker, D. Trystram

Scheduling Problems with Two Competing Agents

Alessandro Agnetis (University of Siena)

We consider the scheduling problems arising when two agents, each with a set of nonpreemptive jobs, compete to perform their respective jobs on a common processing resource. Each agent wants to minimize a certain objective function, which depends on the completion times of its jobs only. The objective functions we consider are maximum of regular functions (associated with each job), number of late jobs, and total weighted completion times. We obtain different scenarios, depending on the objective function of each agent, and on the structure of the processing system (single machine or shop). For each scenario, we address the complexity of various problems, namely, finding the optimal solution for one agent with a constraint on the other agents cost function, finding single Pareto-optimal schedules (i.e., such that a better schedule for one of the two agents necessarily results in a worse schedule for the other agent), and generating all Pareto-optimal schedules. Characterizing the set of Pareto-optimal schedules is also relevant if the situation is modeled as a bargaining game, since a schedule corresponding to a Nash solution always belongs to such schedules.

Joint work of: Agnetis, Alessandro; De Pascale, Gianluca; Mirchandani, Pitu B.; Pacciarelli, Dario; Pacifici, Andrea

Scheduling Algorithms for Transactional Memory

Michael A. Bender (SUNY at Stony Brook)

We present algorithms for scheduling attempts to access memory in shared-memory systems. Our algorithms are based on randomized backoff. Randomized exponential backoff is used in many systems, including Ethernet, satellite networks, email retransmission, optical switching, and transactional memory conflict resolution. Most analysis of randomized exponential backoff assumes some statistical property of the arrivals of jobs, e.g., the arrival rates are governed by a Poisson process or by a self-similar-process. We present our results analyzing worst-case arrivals, where an adversary chooses the times when jobs start trying to run. We give a thorough analysis of a range of backoff strategies for the special case in which all the jobs arrive at the same time. For the case of online arrivals, we assume an adversarial-queuing-theory model and prove upper and lower bounds on the average arrival rates for stability and instability of binary exponential backoff.

Joint work of: Bender, Michael A.; Farach-Colton, Martin; He, Simai; Kuszmaul, Bradley; Leiserson, Charles E.

Complexity of Mean Flow Time Scheduling Problems with Release Dates

Peter Brucker (Universität Osnabrück)

We present a polynomial time algorithm for solving the scheduling problem $P \mid r_i, p_i = p, pmtn \mid \sum C_i$. It is based on a linear program for calculating optimal completion times and a network flow problem for calculating a corresponding schedule. These results are used to derive a polynomial algorithm for the related open shop problem $O \mid r_i, p_{ij} = 1 \mid \sum C_i$ as well. We finally prove unary NP-hardness for the problem $P \mid r_i, pmtn \mid \sum C_i$.

Keywords: Parallel machine scheduling, linear programming, NP-hardness, maximum flow

Joint work of: Brucker, Peter; Kravchenko, Svetlana A.

Aircraft Departure Scheduling at Heathrow Airport

Edmund Burke (University of Nottingham)

We will discuss a methodology that will form the basis of a decision support system to advise the runway controllers at London Heathrow Airport. This work was funded by the UK Engineering and Physical Sciences Research Council (EPSRC) and National Air Traffic Services Ltd through the Smith Institute for Industrial Mathematics and Systems Engineering. Heathrow Airport is one of the busiest airports in the world. We will present the constraints that affect the problem and discuss a range of issues that are critical to the effective automated generation of solutions. We will present four search methodologies and discuss an empirical evolution of their performance on real historical data. This will demonstrate the effectiveness of these techniques when compared to manually produced real world solutions. Finally we will briefly discuss some future research directions.

Joint work of: Burke, Edmund; Atkin, Jason

An Approximation Algorithm for Periodic Real-Time Tasks with Workload-Dependant Running Time

Frank Drews (Ohio University)

This talk addresses the problem of resource allocation for distributed real-time periodic tasks. The objective of resource allocation is to produce an allocation that maximizes the input size (or workload) that applications can handle, and hence to delay possible reallocations as long as possible. Hill climbing, random search, simulated annealing, and first-fit heuristics are presented and evaluated via simulation. As we show here, the first-fit greedy heuristic produces solutions that are often close to optimal. In particular, we show analytically that the first-fit algorithm is guaranteed to produce a solution that is at least $1/2.42$ of optimal, asymptotically, under certain reasonable restrictions on the running times of tasks in the system. Moreover, we show that if at most 12% of the system utilization is consumed by input independent tasks (e.g., constant time tasks), then first fit is guaranteed to produce a solution that is at least $1/3$ of optimal, asymptotically. This suggests that the first-fit algorithm, in combination with other local improvement strategies, may be a reasonable approach for resource allocation in dynamic real-time systems.

Joint work of: Drews, Frank; Juedes, David; Welch, Lonnie; Ecker, Klaus

University Course Timetabling

Andreas Drexl (Universität Kiel)

The timetabling problem consists in scheduling lectures between teachers and students in a given period of time. Timetabling can be broadly classified into three areas, that is, school timetabling, university course timetabling and examination timetabling. Problems differ from each other based on the type of institution involved, viz. university or school. University course timetabling comprises two problems, i.e., lecture to teacher as well as lecture to room and period assignment. This paper addresses the second problem, and it has two main contributions: First, we provide a non-traditional 'teacher-based' mixed-integer program with a huge number of binary variables. Second, we present a column generation algorithm for the approximate solution of large-scale instances which is capable to solve large instances close to optimality in a fair amount of computation time.

Keywords: University course timetabling, huge integer program, column generation, computational results

Joint work of: Drexl, Andreas; Haase, Knut

Generalized Scheduling Problems

Moshe Dror (University of Arizona)

We examine a selective list of combinatorial optimization problems in NP with respect to inapproximability (Arora and Lund, 1997) given that the ground set of elements N has additional characteristics. For each problem in this paper, the set N is expressed explicitly by subsets of N either as a partition or in the form of a cover. The problems examined are generalizations of well known classical graph problems and include the minimal spanning tree, the assignment problem, a number of elementary machine scheduling problems, bin-packing, and the TSP. We conclude that for all these generalized problems the existence of PTAS (polynomial time approximation scheme) is impossible unless P=NP. This suggests a partial characterization for a family of inapproximable problems. For the generalized Euclidean TSP we prove inapproximability even if the subsets are of cardinality two.

Joint work of: Dror, Moshe; Orlin, James B.

(Some) Combinatorics in Divisible Load Scheduling

Maciej Drozdowski (Politechnika Poznanska)

In this presentation we consider computational complexity of divisible load scheduling problem. So far divisible load scheduling problem has been considered as a polynomially solvable one. However, it is the case for a given set of processors to be used, and for a given sequence of activating them. In such a case optimum distribution of the load can be found in polynomial time by use of linear programming. If the set of processors to be used is arbitrary, then this problem is NP-hard when processors have restricted intervals of availability, or computation startup time, or fixed cost charging for being used. The case of non-zero communication startup times, and limited processor speeds is open from the complexity point of view. Special cases of this problem can be solved in pseudopolynomial time. This problem has also some similarities with the half-product minimization problem.

Keywords: Divisible loads, computational complexity

Joint work of: Drozdowski, Maciej; Lawenda, Marcin

Batch Processing with Interval Graph Compatibilities

Gerd Finke (IMAG - Grenoble)

We analyze batch-scheduling problems that arise in connection with certain industrial applications. The models concern processing on a single max-batch machine with the additional feature that the tasks of the same batch have to be compatible. Compatibility is a symmetric binary relation – the compatible pairs are described with an undirected "compatibility graph", which is often an interval graph according to some natural practical conditions that we present. We consider several models with varying batch capacities, processing times or compatibility graphs; we provide new polynomial time algorithms - and some NP-completeness results.

Joint work of: Finke, G.; Jost, V.; Queyranne, V; Sebo, V.

Due Date Assignment and Scheduling via Modular Decomposition

Valery Gordon (Belarus Academy of Sciences - Minsk)

A single machine scheduling problem of minimizing holding costs with no tardy jobs is considered subject to the SLK due date assignment rule. We show that the problem of minimizing non-increasing function of slack and total weighted earliness or total weighted exponential earliness has a polynomial-time solution if precedence constraints are given by a series-parallel graph or by a graph which can be decomposed in such a way that the size of "building" blocks is limited. In the latter case, to solve the problem we use the combination of dynamic programming and modular decomposition. The work is supported in part by INTAS and ISTC.

Joint work of: Gordon, V.; Proth, J.-M.; Strusevic, V.

Solving the Gate Assignment Problem at Amsterdam Airport Schiphol

Han Hoogeveen (Utrecht University)

Deciding which gate each flight at an airport should be assigned to seems like a very simple problem. Unfortunately it is not as simple as it used to be. Besides that the volume of air traffic has increased very much, also the need for satisfying an increasing number of constraints has made solving the gate assignment problem very complex. Based on the constraints that are actually used at Schiphol Airport we have formulated a model and solved this using techniques from Linear Programming. We have used a novel ILP formulation that is not based

on individual flights but on sets of flights that can be assigned to one gate, the so-called gate plans. Instead of solving the ILP with the whole set of feasible gate plans, we have selected a number of gate plans that may be assumed to be 'useful' and have solved the ILP for this much smaller number of gate plans. Here the gate plans were selected by applying column generation, where we added in each iteration a small number of other, promising gate plans as well. Experiments with real-life data show that this approach allows us to generate solutions that are close to optimal in a reasonable amount of time.

Joint work of: Hoogeveen, Han; van den Akker, Marjan; Diepen, Guido; Smeltink, Job

Single Machine Scheduling with Job Values Dependent on Their Completion Times

Adam Janiak (TU Wroclaw)

The work deals with two single machine scheduling problems in which the sum of values of all the jobs is maximized. In the first problem the value of job is characterized by a stepwise non-increasing function with one or more moments at which a change of job value occurs, whereas in the second one the value is described by an exponential non-increasing function. Establishing an order of processing of datagrams which are sent by router is a practical example of application of such problems.

It is proved, that a special case of the first problem – with a single common moment of change of job value and the zero value of jobs after this moment – is NP-hard. Moreover we show that more general case – only with a single moment of change of job value – can be solved in pseudo-polynomial time. In order to prove it, we show that some existing pseudo-polynomial time algorithm for the problem of minimization of the weighted number of late jobs can be adopted to solve our problem.

The second problem of maximization of the total job values is reformulated so that the sum of loss of job values should be minimized. It is proved that the problem is NP-hard.

Both NP-hardness proofs were done on the basis of Partition problem.

Finally, we solve some special cases of both problems in polynomial time and, in order to solve the general versions of the problems, we construct and experimentally test a number of heuristic algorithms.

Joint work of: Janiak, Adam; Krysiak, Tomasz

Single Machine Scheduling with Earliness and Tardiness Costs

Joanna Józefowska (Politechnika Poznanska)

The problem of scheduling jobs on a single machine to minimize the total earliness and tardiness costs is considered. This kind of problems is often met in so-called Just-in-Time environment, where jobs are penalized for being completed before their due dates (storage costs) as well as for being completed after their due-dates (customer dissatisfaction costs). The model with linear cost functions and individual due dates is considered with idle time allowed. The problem has been proved to be NP-complete even for the non-weighted case. However, given a sequence of jobs the vector of optimal starting times can be found in polynomial time. In this presentation we describe an $O(n \log n)$ algorithm to solve this problem which improves the results known in the literature. Moreover, we propose a branch and bound as well as dynamic programming procedure to find optimal sequences of jobs. The branch and bound procedure solves instances with up to 20 jobs. Finally, we propose a tabu search heuristic and evaluate its effectiveness on a set of randomly generated instances. The proposed tabu search algorithm solves instances with up to 1000 jobs within 45 minutes of CPU time. Average relative deviation from optimum (calculated for small instances) does not exceed 4%. Further research is planned to examine related problems with general convex cost functions and non-zero job release dates.

Joint work of: Józefowska, Joanna; Bauman, Jakub; Muras, Micha

New Precedence Theorems for One-Machine Weighted Tardiness

John J. Kanet (University of Dayton)

In an earlier paper by Emmons (1969), the problem of sequencing jobs on a single machine in order to minimize total tardiness was analyzed. Emmons provided three theorems for specifying precedence relations for pairs of jobs. His theorems apply when the tardiness penalty for each job grows at the same rate. Rinnooy Kan, Lageweg, and Lenstra (1975) later extended Emmons's theorems to the case when job tardiness penalties can grow at different rates for different jobs. Provided here is a set of stronger theorems than those of Rinnooy Kan, et al., which more fully exploits the special properties of the weighted tardiness function, allowing for greater reduction of the solution space.

Hyper-heuristcs : An Emerging Search Technology

Graham Kendall (University of Nottingham)

We introduce an emerging methodology in search and optimisation. The aim of this approach, which has been termed hyper-heuristics, is to raise the level of generality at which optimisation systems can operate. The hope is that hyper-heuristics will lead to more general systems that are able to handle a wider range of problem domains than current (meta-)heuristic approaches, which tend to be customised to a particular class of problems or even specific problem instances. Hyper-heuristics are broadly concerned with intelligently choosing the right heuristic (or algorithm) at each decision point to produce solutions which are good enough, soon enough, cheap enough. In this talk we will give a brief history of this emerging area and discuss the progress made to date.

Scheduling with Variable Intensity Activities and Feeding Precedence Constraints

Tamás Kis (Hungarian Academy of Sciences)

In this problem a set of variable intensity activities must be scheduled on a set of continuously divisible renewable resources while respecting release times, deadlines, feeding precedence constraints (explained below), maximum intensity, and resource constraints, the objective being to minimize the cost of using extra resource capacities. A "feeding" precedence constraint between a pair of activities i and j with parameter p prescribes that (1) p percent of activity i must be completed before activity j may start, and (2) $f(i,t) \geq f(j,t)$ at any time period t, where $f(i,t)$ and $f(j,t)$ denote the fraction of activity i and j, respectively, completed up to time t. We formalize this problem as a mixed integer-linear program and present various polyhedral results. The effectiveness of the cuts derived are demonstrated by computational results obtained by a branch-and-cut algorithm. Such problems occur in production planning in discrete manufacturing.

Keywords: Production planning, mixed integer-linear programming, cutting planes, network flows

Scheduling Chains with Identical Jobs and Constant Delays on a Single Machine

Sigrid Knust (Universität Osnabrück)

We study the single-machine problem $1|chains(l), p_j = p|\sum C_j$ in which jobs with constant processing times and generalized precedence constraints in form of chains with constant delays are given. One has to schedule the jobs on a single machine such that all delays between consecutive jobs in a chain are satisfied and the sum of all completion times of the jobs is minimized. We show that this problem is polynomially solvable.

Keywords: Scheduling, time-lags, delays, complexity results

Joint work of: Brucker, Peter; Knust, Sigrid; Oğuz, Ceyda

Scheduling of Work and Rework Processes

Mikhail Kovalyov (Belarus Academy of Sciences - Minsk)

Two problems of reverse logistics are studied. In the first problem, there is one disassembly plant and F remanufacturing plants. Starting times for all the plants are given. The disassembly plant produces parts of F different types and parts of type f are delivered to and are used by the remanufacturing plant f, $f = 1, ..., F$. A setup occurs if the disassembly plant switches from one part type to another part type. The objective is to find a batch sequence of parts at the disassembly plant such that the total setup cost is minimized and each remanufacturing plant keeps its production continuous. The problem is proved to be NP-hard. A dynamic programming algorithm is derived for the general case. For $F = 2$, the problem is solved in a polynomial time. For unit setup costs, an approximation algorithm is derived, which produces a solution with the number of setups at most F times the minimum number. In the second problem, the same facility is used to manufacture items of the same product and to remanufacture defective items. Demands for good quality items are given. The percentage of defective items is assumed to be known. Furthermore, they are evenly distributed among good quality items. Items are produced in batches. Batch processing includes two stages. In the first stage, all items of the batch are manufactured. In the second stage, defective items of the same batch are remanufactured. Some of them can be disposed of. Each stage is preceded by a setup time. While the manufacturing time is the same for all items, the remanufacturing time for a defective item depends on its holding time between the two stages. The objective is to find batch sizes for a schedule with a given number of good quality items such that the sum of batching, holding, shortage and disposal costs is minimized. Properties of an optimal solution are established and efficient dynamic programming algorithms are developed. This research was supported in part by INTAS.

Chain Constrained Scheduling in Survivability Applications

Axel Krings (University of Idaho)

This research addresses problem transformations of two survivability applications, i.e. (1) survivable migratory agents and (2) network backup services, into scheduling problems with chain precedence. In the first application the theoretical implications of determining agent traversal routes in a secret sharing scheme are considered. A transformation is shown that utilizes strong and weak chain precedence in order to optimally assemble agents carrying different shares at

processors. The second application is transformed into scheduling models using chain precedence for backup tasks. Infrastructure specific parameters, e.g. link bandwidth, are used to introduce chain precedence on tasks, thereby avoiding fan-outs that could lead to scheduling infeasibility. The application transformations allow for solving the original problem in the transformation domain.

Keywords: Scheduling, chain precedence, problem transformation, survivability application

Batch and Cyclic Scheduling of 1- and 2-processor Unit Time Tasks on Dedicated Processors

Marek Kubale (TU Gdansk)

This talk consists of two parts. In the first part we consider a dedicated model of batch scheduling of unitary 1- and 2-processor tasks. We show that there exists a polynomial-time algorithm based on graph (edge) coloring for scheduling such tasks which guarantees a solution of length at most one unit of time more than the optimal one. In the second part we analyze a dedicated model of cyclic scheduling of unitary 1- and 2-processor tasks. We show that such a cyclic schedule is shorter than the corresponding batch schedule by less than a unit of time. Finally, we show that there exists a polynomial-time 1.25-approximation algorithm for scheduling such tasks.

Preemptive Open Shop Scheduling with Multiprocessors

Wieslaw Kubiak (University of Newfoundland)

This presentation addresses a multiprocessor generalization of the preemptive open-shop scheduling problem. The set of processors is partitioned into two groups and the operations of the jobs may require either single processors in either group or simultaneously all processors from the same group. We consider two variants depending on whether preemptions are allowed at any fractional time point or only at integral time points. We shall show that the former problem can be solved in polynomial time, and provide sufficient conditions under which the latter problem is tractable. Applications to course scheduling and hypergraph edge coloring are also discussed.

Joint work of: Kis, Tamás; Kubiak, Wieslaw; de Werra, Dominique

Approximation Algorithms for the Packing of Multibin Objects

Pierre Lemaire (IMAG - Grenoble)

We consider the problem of packing multibin objects. Formally, an object is defined by its height (the space it occupies in each bin it is packed into) and its width (the number of parts it is made of). Each part of an object must be packed into a different bin. The bins are identical. For a given solution, the height of a bin is the sum of the heights of the objects packed into it. The objective is to minimize the height of the highest bin. This problem is strongly connected to scheduling problems: it is a relaxation of multiprocessor-task scheduling and an extension of parallel machine scheduling. It is NP-hard but solvable in pseudo-polynomial time for a fixed number of bins. In this case, it admits a PTAS and a FPTAS. However, these methods are useless from a practical point of view. We are interested, here, in fast and efficient approximation algorithms. A first and very simple heuristic is Next-Bins (NB): the first object is packed into the first bins, the next object in the next bins and so on. This algorithm is optimal if the objects are all of the same height. In the general case, it may produce really bad solutions; if the objects are sorted by non-increasing heights, however, it is a 2-approximation. A second algorithm is Optimist-Fit (OF): a reference value is computed (e.g. a good lower bound, such as the average volume per bin); then the objects are packed into the first bins where they fit under this limit. If no such bins exist, the object is packed into the lowest bins. An extension of this algorithm (OF-X) consists in performing the above procedure for every pertinent reference value (i.e. within a lower and a upper bound). The exact guarantee of this algorithm is not known, but it is conjecture not to exceed the optimal value plus the height of the highest object, if these objects are sorted by non-increasing widths. A third algorithm (DC) is based on the divide and conquer principle. Partial solutions are made, initially one per object. Then the partial solutions are iteratively merged by two, the lowest bins of one with the highest bins of the other. This algorithm is a 2-approximation (indeed, the solutions produced never exceed the optimal value plus the height of the highest object), but better guarantees are expected from non-blind merging procedures. A forth algorithm is Best-Fit (BF): an object is always packed into the lowest bins. The on-line version is a 2-approximation. The decreasing version (DBF), for which the objects are sorted by non-increasing heights, is a 4/3-approximation. This algorithm may take into account some pre-assignments of objects. Experiments on random instances show that these algorithms perform very well (however NB has no interest, with regards to the others). An interesting feature of these algorithms (in particular DBF, OF-X and DC) is that they perform well each on different kinds of instances, and therefore are very complementary. Indeed, when one performs poorly, another one performs very well. This explains that more than 75% of the instances are solved optimally, and that the average gap to the optimal solution is less than 0.3%. For future work, the OF and DC approaches should be developed, inspired respectively by the multifit and the

differencing method for parallel machine scheduling. Exact methods should also be investigated.

Joint work of: Lemaire, P.; Finke, G.; Brauner, N.

Bicriteria Moldable Task Scheduling

Gregory Mounie (INRIA Rhône-Alpes)

The Moldable Tasks model is promising for scheduling parallel tasks on parallel computers. Hall et al (1984) designed a general scheme for independent tasks optimizing the C_{max} and $\sum(w_i C_i)$ criteria, assuming there is an approximation algorithm for C_{max}. For monotonic moldable tasks, a 2-shelf scheme gives a 3/2 + epsilon approximation for C_{max}. This algorithm was adapted into the general scheme to yield a (3,6)-offline approximation algorithm.

Such a sophisticated algorithm is difficult to implement in an actual system because of additional complex constraints. It was therefore simplified such that it can fit into a classic FIFO + backfilling environment. Simulation results show that this simpler version has good average performance.

Joint work of: Mounie, Gregory; Dutot, Pierre-François; Eyraud, Lionel; Trystram, Denis

A Decomposition Method for Batch Production Scheduling

Klaus Neumann (Universität Karlsruhe)

The batch scheduling problem in process industries consists in scheduling batches on processing units such a regular objective function (e.g., the makespan or maximum lateness) is minimized. In addition to processing units with limited capacity, scarce resources such as manpower and storage facilities for intermediate products have to be observed. We model the batch scheduling problem as a resource-constrained project scheduling problem with temporal, renewable-resource, and storage-resource constraints. The project scheduling problem can (approximately) be solved by branch-and-bound methods, truncated branch-and-bound techniques (e.g., filtered beam search), or priority-rule methods. In process industries, there are problem instances with several thousand operations. To deal with such large instances, we present a decomposition approach, where the original problem is decomposed into subproblems with respect to operations such that bottleneck resources are evenly loaded. The subproblems, whose activity sets result from solving an integer linear program, are constructed iteratively and solved by filtered beam search. To obtain a (feasible) solution of the original problem, the individual subschedules are linked together at specific concatenation times.

Berth Allocation as a Moldable Task Scheduling Problem

Ceyda Oğuz (Hong Kong Polytechnic University)

In this paper, the allocation problem of berths to the incoming ships is modeled by moldable tasks scheduling problem. This model considers the tasks as the ships and the processors as quay cranes assigned to these ships. Since the duration of berthing for a ship depends on the number of quay cranes allocated to the ship, the use of moldable task scheduling model is substantiated. In the model, the processing speed of a task is considered to be a non-linear function of the number of processors allocated to it. A suboptimal algorithm, which starts from the continuous version of the problem (i.e. where the tasks may require a fractional part of the resources) to obtain a feasible solution to the discrete version of the problem, is presented. The computational experiments conducted showed that the suboptimal algorithm has a very good average behavior.

Keywords: Berth scheduling, moldable task, resource allocation, parallel computing

Joint work of: Błażewicz, Jacek; Cheng, T.C.E.; Machowiak, Maciej; Oğuz, Ceyda

Scheduling (Planning) for an Industrial Production Facility

Cynthia A. Phillips (Sandia National Labs - Albuquerque)

We consider a planning problem for the production facility responsible for evaluation of the US nuclear stockpile. Weapons arrive from the field, are disassembled, run through a series of custom tests, reassembled, and returned to the field. The facility must schedule jobs consisting of multiple (tree) precedence-constrained tasks. Each task must be completed in a given time window (release date, deadline) in an appropriate facility with an appropriately-sized team of qualified technicians. We wish to plan for a year or more. A plan is a schedule that meets all the above constraints except that facility and technician constraints need only be met at a weekly granularity. Even with this relaxation, future pending work requires more resources than is currently available. Therefore, we wish to compute a plan that requires the fewest additional (facility and technician) resources to meet all the constraints. The associated mixed-integer program (MIP) has hundreds of thousands of variables and cannot be solved with commercial serial MIP solvers. We experimentally investigate alpha-point heuristics for this problem. Alpha-point heuristics were originally designed for cleaner problems, where they can give provably good approximations for minimizing average weighted completion time. We showed that alpha-point heuristics can reasonably approximate an optimal plan, usually within 10-15% of optimal, on synthetic data sets designed to mimic the structure of Pantex problems. These heuristics run in minutes for problems that required days or weeks to run to

optimality using serial MIP programs. This is the first empirical study to apply alpha-point methods to a problem with hard deadlines. This constraint and the nature of the objective function destroy the structure usually used to prove near-optimality of alpha-point schedules. This shows the global structure of the alpha-point ordering could be stronger than previous known. In our experiments, we also found that using an alpha-point heuristic with the best alpha dominates using alpha fixed to .5 followed by local improvement for that heuristic. Finally, we found that start-based alpha points perform better than completion-based alpha points for hard-deadline problems, probably because they bias for earlier starts for long jobs. We are aware of no other studies that compare the two types of alpha points.

Rescheduling

Chris N. Potts (University of Southampton)

This paper considers scheduling problems where the processing of a set of jobs has been scheduled (i.e., planned) to minimize a classical cost objective, under the assumption that those jobs are all available at the start of the planning horizon. Before processing starts, however, the availability of a subset of the jobs is delayed. Therefore, the decision maker needs to adjust the existing schedule to allow for the initial unavailability of these jobs, but without causing excessive disruption to the schedule and expensive resource reallocations. The limit on allowable disruption is measured by the maximum time disruption to any job, between the original and adjusted schedules. For four classical scheduling objectives, we provide both a computationally efficient optimal algorithm and an intractability proof showing that such an algorithm is the best possible type of result. Also, for each problem considered, we provide an efficient approximate solution procedure, show that its worst-case performance ratio is a small attainable constant, and demonstrate computationally that its average performance is very close to optimal. Finally, we provide for each problem a fully polynomial time approximation scheme. Our work refocuses the extensive literature on scheduling problems towards rescheduling issues which are important because of the frequency with which disruptions occur in manufacturing practice.

Joint work of: Potts, Chris N.; Hall, Nicholas G.

Minimizing Travels by Maximizing Breaks in Round Robin Tournament Schedules

Celso Carneiro Ribeiro (University of Rio de Janeiro)

We investigate the relation between two aspects of round robin tournament scheduling problems: breaks and distances. We show that the breaks maximization problem and the distance minimization problem are equivalent for a new class

of uniform problem instances. This relation is used to prove the optimality of solutions found by a heuristic to the mirrored traveling tournament problem. The uniform instances solved to optimality are the largest solved to date and shed some light on the complexity of this problem.

Joint work of: Ribeiro, Celso Carneiro; Urrutia, Sebastian

Master Scheduling in Make-to-order Discrete Manufacturing

Tadeusz Sawik (AGH - Krakow)

IP formulation with various cutting constraints is proposed for master scheduling in make-to-order discrete manufacturing. The objective is to allocate customer orders among planning periods to minimize number of tardy and early orders as well as to balance machine assignments over a planning horizon. Numerical examples modeled after a real-world long-term scheduling in the electronics industry are provided.

Keywords: Production scheduling, Hybrid flowshop, Due-date scheduling. Integer programming.

Partitioning a Permutation Graph: Exact Algorithms and an Application

Frits C.R. Spieksma (University of Leuven)

We discuss the problem of partitioning a permutation graph into cliques of bounded size. We describe a real-life application of this problem encountered at a manufacturing company in the Netherlands and we present two exact algorithms for solving this problem. The first algorithm is a branch-and-price algorithm, based on an integer programming formulation. The pricing problem can be formulated as a longest path problem and can be solved efficiently by dynamic programming. The second algorithm is an enumeration algorithm based on the concept of bounded clique-width. This algorithm was motivated by a special structure that is present in the real-life instances that were used for computational experiments. Both algorithms are tested on a number of real-life instances and randomly generated instances. From the computational results we conclude that both the real-life and the random instances can be solved satisfactorily by the branch-and-price algorithm. The enumeration algorithm performs really well in case of the real-life instances (99% of the instances are solved within a second), but the random instances cannot be solved efficiently, due to the large number of different lengths in the input.

Joint work of: Spieksma, Frits C.R.; Moonen, Linda

Solution Methods for Two-Machine Flow Shop with Weighted Late Work Criterion

Malgorzata Sterna (Politechnika Poznanska)

In the work, we compared three different solution methods for the two-machine flow shop scheduling problem with the common due date and the weighted late work criterion, which is known to be binary NP-hard. We implemented and tested the dynamic programming approach, the enumerative method and the list scheduling algorithm with different static and dynamic priority dispatching rules.

Computational experiments showed that intuitive assumptions on time efficiency of different solution methods may not be always confirmed in practice. Each algorithm has its advantages and none of them can be rejected from the consideration a priori. During the tests the enumerative method appeared to be more efficient than the dynamic programming one owing to some special features of the problem under consideration. Moreover, computational experiments supported the general opinion on high utility of the list scheduling algorithm. It constructed solutions of a good quality in the reasonable amount of time.

Keywords: Flow shop, late work, dynamic programming, enumerative method, list scheduling

Joint work of: Błażewicz, Jacek; Pesch, Erwin; Sterna, Malgorzata; Werner, Frank

Planning Machine Maintenance in Two-Machine Shop Scheduling

Vitaly Strusevich (The University of Greenwich)

We consider the two-machine open shop and two-machine flow shop scheduling problems in which there is a mandatory maintenance period on each machine $L \in \{A, B\}$. If the maintenance period (MP) on machine L starts at time t, then its length is equal to $\Delta_L(t) = \alpha_L + f_L(t)$, where α_L is a given positive machine-dependent constant, and $f_L(t)$ is a given monotone non-decreasing machine-dependent function such that $f_L(0) = 0$, $L \in \{A, B\}$. During an MP no job can be processed on the corresponding machine. Both MPs can be run in parallel. In our model the length of an MP depends on its start time, which is quite reasonable, since a machine comes to an earlier maintenance in a better condition, so that less time is needed for fixing it. The makespan is defined as the maximum completion time of all jobs and all MPs.

We completely resolve complexity and approximability issues of these problems. The open shop problem is shown to be solvable in linear time by an extension of the algorithm by Posner and Lu for the problem with a machine

release date. By contrast, the flow shop problem is proved binary NP-hard and pseudopolynomially solvable by dynamic programming if only one of the machines is subject to maintenance and the length of the MP depends linearly on its start time t. We also present a fully polynomial approximation scheme and a fast 3/2-approximation algorithm for the latter problem.

Joint work of: Strusevich, Vitaly; Kubzin, M.A.

On-Line Scheduling of Multiprocessor Jobs with Idle Regulation

Andrei Tchernykh (CICESE Research Center - San Diego)

In this paper, we focus on on-line scheduling of multiprocessor jobs with emphasis on the regulation of idle periods in the frame of general list policies. We considered non-clairvoyant scheduler that has no information about the jobs other than the number of unfinished jobs in the system and their processor requirements. We proposed to solve on-line scheduling problems using batch scheduling under a generic framework of two-phases of list scheduling which successively combine sequential and parallel executions of the jobs. Scheduling strategies used a dynamic idle regulation approach in order to improve system behavior has receive attention recently. In this paper, we showed that tuning of the parameter that regulates idle times gives a possibility to find a tradeoff between avoiding idle processors by starting parallelization sooner when more tasks are parallelized and delaying their parallelization until a smaller number of jobs is available. A tradeoff is a minimum of the performance guarantee for a parallel computer system with fixed number of processors and known overhead. The minimum depends on the penalty factor of the system and applications, and can be adapted to the change of quality of the workload on-line, based on runtime measuring of workload characteristics. The main result is to demonstrate that it is possible to estimate the amount of resources that should remain idle for a better regulation of the load and to obtain approximation bounds.

Full Paper: A.Tchernykh, D.Trystram. On-Line Scheduling of Multiprocessor Jobs with Idle Regulation. Parallel Processing and Applied Mathematics, Wyrzykowski et al. (Eds.): PPAM 2003, LNCS 3019, pp. 131-144, 2004. Springer-Verlag. Berlin, Heidelberg

Joint work of: Tchernykh, Andrei; Trystram, Denis

A New Approach Based on Caches for Scheduling

Denis Trystram (INRIA Rhône-Alpes)

In this talk, we investigate a new approach for improving the results of scheduling algorithms in the context of parallel computations on Internet. Basically, we address the problem of scheduling efficiently the tasks of programs corresponding to series of request submitted to a web server that is connected to a parallel cluster. These computations are decomposed into computational units that may have some redundancies. We proposed a modelization of this general problem using graphs where some nodes are redundant. Such nodes may be put into caches in order to avoid their computation. The problem is polynomial where the cache has no limitation, but it is NP-hard in the case of limited size. In this case, we propose an algorithm that minimizes the user response time. This approach has been tested on an actual problem of dynamic requests on geographic maps server in the frame of an European project. Some preliminary results are discussed and open questions are given.

Joint work of: Trystram, Denis; Edi, Euloge; Vincent, Jean-Marc

Load-balancing Iterative Computations on Distributed Heterogeneous Platforms (the Impact of the Network Model)

Frédéric Vivien (ENS - Lyon)

This work is devoted to mapping iterative algorithms onto heterogeneous clusters. The application data is partitioned over the processors, which are arranged along a virtual ring. At each iteration, independent calculations are carried out in parallel, and some communications take place between consecutive processors in the ring. The question is to determine how to slice the application data into chunks, and to assign these chunks to the processors, so that the total execution time is minimized. One major difficulty is to embed a processor ring into a network that typically is not fully connected, so that some communication links have to be shared by several processor pairs. We focus on the impact of the network model on the complexity of the problem.

Joint work of: Renard, Hélène; Robert, Yves; Vivien, Frédéric

Fluid Approach to Scheduling: Continuous Linear Programs, Processing Networks with Infinite Virtual Queues, and Maximum Pressure Policies

Gideon Weiss (Haifa University)

We present an approach to control systems with many items evolving in a network via fluid approximation. This approach is asymptotically optimal for many problems including the minimization of holding costs in a job shop with a fixed set of routes and many jobs on each route. The fluid approach consists of

- Formulate fluid model of the system
- Optimize fluid model by solving SCLP (separated continuous linear program) with simplex algorithm
- Translate fluid solution to a policy for the discrete system

The fluid solution partitions the time horizon to intervals. In each interval there is a set of non-empty buffers with flows through empty buffers. This is translated into a processing network with infinite virtual queues (at the non-empty buffers) each of which has a nominal flow rate and standard queues (at the empty buffers). Maximum pressure policy as defined by Dai & Lin is then used to maintain the nominal flow rates while keeping the standard queues path-wise stable.

Scheduling with Common Due Date, Earliness and Tardiness Penalties for Multi-Machine Problems: A Survey

Frank Werner (Universität Magdeburg)

In the talk we review some recent results in the area of scheduling jobs with common due date on a set of machines, where the nonregular optimization criterion depends on the earliness and tardiness penalties of the jobs. The focus in the talk is on shop scheduling problems. In this case, the objective functions are either similar to the single-machine problems, or additionally intermediate storage costs between the operations of a job have to be taken into account. We discuss complexity issues, exact and heuristic algorithms. For two-machine flow shop problems detailed computational results with an enumerative algorithm and several heuristics are presented.

Keywords: Common due date scheduling, nonregular criteria, earliness/tardiness penalties

Joint work of: Werner, Frank; Lauff, Volker

Agents in Rostering Systems

Patrick de Causmaecker (KaHo St.-Lieven - Gent)

Timetabling is investigated as a detailed planning problem at distinct levels in a hierarchical organisation. In the process of building a system for such a problem, we identified the need for a universal framework in which low level rostering problems as well as higher level negotiation issues could be handled. The framework should at all times be able to cope with incomplete information and limited connectivity. The reasons for this are organisational as well as technical. We describe an agent based formulation for the elementary rostering problem. A random algorithm for handling this problem has been developed. Similarities with other heuristics are apparent. This formulation and algorithm are then generalised into a negotiation scheme that can be used at all levels of the hierarchy. This approach leads to some questions to be investigated in future research.

Keywords: Timetabling, Agents, Distributed

Some Combinatorial Problems of Authentication in Ad Hoc Mobile Networks

Dominique de Werra (EPFL - Lausanne)

For securing communications in a mobile network, users have to organize an authentication system based on public-key cryptography. Each user is represented by a node v of a graph $G = (V, U)$ and there is an arc (v, w) if v has a certificate of authentication for w; authentication of w by v can also occur through a chain of certificates, i.e., through a path from v to w in G. The problem of authentication is formulated as follows: given a strongly connected graph $G = (V, U)$, one has to assign to each node v a partial subgraph $G_v = (V_v, U_v)$ (certificates of authentication stored by user v) in such a way that for any users v, w the union of graphs G_v and G_w (obtained by taking the arcs of G_v and those of G_w) contains a path from v to w and a path from w to v. In order to simplify the authentication procedures the collection of graphs G_v must be such that $max_v|U_v|$ is minimum under the additional requirement that each node v should not occur in more than u_o of the subgraphs G_v (u_o is a given integer). The complexity of the associated decision problem is apparently still open. Bounds on $max_v|U_v|$ are given and some special cases are studied, like complete graphs and tree-like graphs. In addition some heuristic procedures are sketched and some computational results are presented.

Keywords: Authentication, strongly connected graph, complete graphs, certificate graph

Joint work of: de Werra, Dominique; Jaccard, T.; Schindl, D.