DAGSTUHL-SEMINAR
SCHEDULING IN COMPUTER AND MANUFACTURING SYSTEMS
June 2-7, 2002

SCHEDULE

Monday - June 3, 2002

Session 9:00-12:00 Chairman: Jan Weglarz

9:00-9:30 Dominique de Werra "Weighted coloring and batch scheduling"
9:30-10:00 Wieslaw Kubiak "On Fraenkel's and Small Deviations Conjectures"
10:00-10:30 Klaus Neumann "Batch scheduling in process industries: case study and solution method"
11:00-11:30 Adam Janiak "Parallel Machine Scheduling Problems with Due Interval Assignment"
11:30-12:00 Maciej Drozdowski "Performance limits of divisible load processing under limited communication buffers"

Session 16:00-18:00 Chairman: Ed Coffman

16:00-16:30 Joe Leung "Minimizing total completion time on parallel machines with deadline constraints"
16:30-17:00 Jacek Blazewicz "On the relation between different graph representations of precedence constraints"
17:00-17:30 Axel Krings "The impact of hybrid fault models on scheduling for survivability"
17:30-18:00 Malgorzata Sterna "Production optimization in FMS environment"

Tuesday - June 4, 2002

Session 9:00-12:00 Chairman: Joe Leung

9:00-9:30 Joanna Jozefowska "On discrete-continuous scheduling problems with the mean flow time minimization"
9:30-10:00 V. Sridharan "Job due date setting in manufacturing"
10:00-10:30 Valery Gordon "Single machine scheduling with due dates depending on processing times"
11:10-11:30 Peter Brucker "Complexity results for shop problems with transportation delays"
11:30-12:00 Marek Kubale "Cylindrical chromatic scheduling of dedicated 2-processor tasks"

Session 16:00-18:00 Chairman: Klaus Neumann

16:00-16:30 Jose Nino-Mora "Scheduling and routing in queuing systems: restless bandits index policies"
16:30-17:00 Erwin Pesch "Airport gate management"
17:00-17:30 Gideon Weiss "Scheduling manufacturing systems via fluid model optimization by solving continuous linear programs"

17:30-18:00 Hans Kellerer “Semi on-line multiprocessor scheduling with given total processing time”

Wednesday - June 5, 2002

Session 9:00-12:00 Chairman: Dominique de Werra

9:00-9:30 Mikhail Kovalyov “Single supplier scheduling for multiple deliveries”
9:30-10:00 Marino Widmer “Designing a timetable under strict legal constraints “
10:00-10:30 Frank Padberg “A stochastic scheduling model for software projects”
11:00-11:30 Cynthia A. Phillips “Real-time node allocation on commodity-based supercomputers” Part I
11:30-12:00 Michael Bender “Real-time node allocation on commodity-based supercomputers” Part II

Thursday - June 6, 2002

Session 9:00-12:00 Chairman: Peter Brucker

9:00-9:30 Olivier Braun “Parallel processor scheduling with limited number of preemptions”
9:30-10:00 Gerd Finke “Relaxing multiprocessor scheduling to multibin packing problems”
10:00-10:30 Jurgen Zimmermann “Scheduling of investment projects in manufacturing systems”
11:00-11:30 Frank Werner “Unit-time job shop scheduling via mixed graph coloring”
11:30-12:00 Han Hoogeven “Accept or reject? The odds are…”

Session 16:00-18:00 Chairman: Erwin Pesch

16:00-16:30 Jay Sethuraman “Scheduling & routing in adversarial queuing networks”
16:30-17:00 Nadia Brauner “Complexity results in a just in time scheduling problem”
17:00-17:30 Denis Trystram “A new clustering algorithm for scheduling with large communication delays”
17:30-18:00 Pierre Dutot “Scheduling divisible tasks on heterogeneous processors”

Friday - June 7, 2002

Session 9:00-12:00 Chairman: Gerd Finke

9:00-9:30 Tadeusz Sawik “Mixed integer programs for scheduling printed wiring board assembly in surface mount technology lines”
9:30-10:00 Artur Czumaj “Selfish traffic allocations”
10:00-10:30 Alix Munier Kordon “A buffer minimization problem for the design of embedded systems”
11:00-11:30 Frank Drews “Incorporating multitasking into a static scheduling model”
11:30-12:00 Antoine Jouglet “A branch-and-bound procedure to minimize total tardiness on one machine with arbitrary release dates”

ABSTRACTS

Weighted Coloring and Batch Scheduling

Dominique de Werra
(Ecole Polytechnique Fédérale de Lausanne, Switzerland)
M. Demange, J. Monnot, V. Paschos
(Université de Paris, France)

For dealing with some batch scheduling problems with incompatibility constraints among jobs, the following graph model has been introduced :
In a graph G = (V, E) where each node v has a weight w(v), a hypostable set S is a subset of V which induces a collection of independent cliques (no edge between two cliques in *S*). A k-hypocoloring is a partition of V into k hypostable sets $S_1,\ldots,S_k$. The weight of a clique K in a set S is the sum of the weights of the nodes of K and w(S) is the maximum of the weights of the cliques K in S. One wants to find for G a value k and a k-hypocoloring $\mathbf{C} = (S_1,\ldots,S_k)$ such that $C_{max}(\mathbf{C}) = \sum\left(w(S_i)\,\middle|\, i = 1,\ldots k\right)$ is minimum.
We show that there is always a solution with $k \leq \Delta$ where Δ is the maximum degree in G. A polynomial algorithm is given where $\Delta = 2$. (The problem is difficult when $\Delta \geq 4$). Finally an enumeration algorithm COCA (connect or contract algorithm) is given for triangle-free graphs. Its "light" version is the classical contraction-connection algorithm for usual colorings of graphs.

On Fraenkel's and Small Deviations Conjectures

Wieslaw Kubiak
(Faculty of Business Administration, Memorial University of New Foundland, Canada)

Brauner and Crama conjectured in 2001 that the only n-product, n > 2, standard instance of the maximum deviation just-in-time sequencing problem with maximum deviation less than 1/2 is made up of the first n nonnegative powers of 2. We prove that this conjecture holds true. Consequently, the answer to the question whether the well known Fraenkal's conjecture holds or not hinges entirely on the sequences' rests, which we also show in this paper.

Batch Scheduling in Process Industries: Case Study and Solution Method

Klaus Neumann
(University of Karlsruhe, Germany)

We discuss the so-called batch scheduling problem in process industries, which consists of scheduling batches on processing units (e.g. reactors, dryers, filters, heaters, agitators) such that the makespan is minimized. First, we present a case study which contains most of the features typical of batch production in process industries: tasks that can be carried out on alternative processing units of limited capacity, sequence-dependent cleaning times of processing units, manpower with reduced size during night shifts and weekends, production breaks, interruptible and non-interruptible tasks, and storage facilities of limited capacity where non-perishable intermediate products are stocked. Second, we show how to model the batch scheduling problem as a multi-mode resource-constrained project scheduling problem. Third, we sketch how to solve the latter problem using a truncated branch-and-bound procedure.

Parallel Machine Scheduling Problems with Due Interval Assignment

Adam Janiak
(Institute of Engineering Cybernetics, Wroclaw University of Technology, Wroclaw, Poland)

The paper deals with parallel identical processors scheduling problems, in which a due interval should be assigned to each job. Due interval is a generalization of well known classical due date and describes a time interval in which job should be finished. Some single processor cases were also considered.
Generally, the problem is to find a sequence of jobs on a single processor or parallel identical processors and an assignment of due intervals to each job, which minimize the considered criterion. We will minimize the following two types of criteria:

1. the maximum of the following weighted three parts: the maximum tardiness, the maximum earliness and the due interval parameters,
2. the sum of the following weighted three parts: the total tardiness, the total earliness and the due interval parameters.

For the single processor cases, some original properties of the optimal solutions have been proved. Due to them, the polynomial time algorithms have been constructed. For the general parallel processors situation, the NP-hardness has been proved even for two processors cases. For the general problems, some approximation algorithms have been constructed and their worst case analyzes have been presented. For some special cases fully polynomial time algorithms have been constructed.

Performance Limits of Divisible Load Processing under Limited Communication Buffers

Maciej Drozdowski and Pawel Wolniewicz
(Institute of Computing Science, Poznan University of Technology, Poznan, Poland)

In this presentation we examined the influence of limited size of communication buffer on the efficiency of processing divisible loads. Divisible loads are computations which can be divided into parts and these parts can be processed in parallel on distributed computers. To finish processing in the shortest possible time an optimum distribution of the load must be calculated. The method of determining load distribution takes into account not only computing speed, but also interconnection system topology, communication medium speed and startup time. In this work we include one more parameter: communication buffer size. We study three archetypal interconnection topologies: stars, ordinary trees, and binomial trees. The problem of finding the optimum load distribution has been reduced to a linear programming model. The results of modeling the performance of parallel systems show that buffer size has quantitative impact on the performance. Limited buffer size causes e.g. causing message fragmentation, or load imbalance. In trees it effectively overturns optimizations proposed in the literature. Changing communication strategies is one of possible ways to overcome this limitation.

Minimizing Total Completion Time on Parallel Machines with Deadline Constraints

Joseph Leung
(Department of Computer Science New Jersey Institute of Technology Newark, NJ 07102, USA)
Michael Pinedo
(Stern School of Business New York University New York, NY 10012, USA)

Consider n independent jobs and m identical machines in parallel. Job j has a processing time p_j and a deadline d_j. It must complete its processing before or at its deadline. All jobs are available at time t = 0 and preemptions are allowed. A set of jobs is said to be feasible if there exists a schedule that meets all

the deadlines; such a schedule is called a feasible schedule. Given a feasible set of jobs, our goal is to find a schedule that minimizes the total completion time. Lawler (1982) asked whether this problem is easy or NP-hard. In this talk we present a polynomial-time algorithm for every m and we show that the general problem with m unrelated machines is strongly NP-hard.

On the Relation between Different Representations of Precedence Constraints

Jacek Blazewicz
(Institute of Computing Science, Poznan University of Technology, Poznan, Poland)
Daniel Kobler
(TmBioscience, Toronto, Canada)

Precedence constraints are a part of a definition of any scheduling problem. After recalling, in precise graph theoretical terms, the relations between task-on-arc and task-on-node representations, we show the equivalence of two distinct results for scheduling problems, i.e. Pm | pmtn, uniconnected activity network | C_{max} = Pm | pmtn, interval order | C_{max}. Furthermore, again using these links between representations, we exhibit several new polynomial cases for various problems of scheduling preemptable tasks on unrelated parallel machines under arbitrary resource constraints.

The Impact of Hybrid Fault Models on Scheduling for Survivability

Axel Krings and Scott Harrison
(Computer Science Department University of Idaho, USA)
Azad Azadmanesh
(Computer Science Department University of Nebraska at Omaha, USA)
Miles McQueen
(Idaho National Environmental and Engineering Laboratories - INEEL, USA)

This research addresses scheduling issues in networked computer systems with survivability requirements, i.e. systems in which essential services must survive malicious acts. In order to achieve survivability based on spatial redundancy, agreement algorithms are needed as a mechanism to consolidate results of individual replicas. The potentially enormous overhead associated with communication and voting schemes of the algorithms put unique burdens on the scheduler as the efficiency of scheduling determines the suitability of the agreement algorithm. We derive agreement task graphs, representing computations, and inter-process communication based on phantom tasks. Task graph primitives are identified and it is shown how their scheduling directly influences the performance of the agreement algorithm. Finally, the notion of dynamic k-of-N precedence is introduced and its impact on slack-time reclaiming for early stopping algorithms is discussed.

Production Optimization in FMS Environment

Jacek Blazewicz and Malgorzata Sterna
(Institute of Computing Science, Poznan University of Technology, Poznan, Poland)
Erwin Pesch
(University of Siegen, Siegen, Germany)

The research concerns the real flexible manufacturing system located at the Poznan University of Technology. It consists of three CNC machines: a lathe, a milling and a measurement machine, a robot with linear routing and four stores. The system works in one shift manner realizing orders for small

series of products. We constructed the formal model of FMS in dedicated machine environment with sequence dependent set-up times. The technological processes are described as the extended mixed shop, i.e. the job shop with open shop sections that reflects the feasibility of the production. We designed the list scheduling algorithm with specialized selection rules for determining a feasible solution of the problem taking into account all system constraints as the participation of an external co-operator in technological processes, limited buffer capacity or dividing the schedule into shifts. Moreover, we proposed three methods of an instance transformation, which allow us to construct different feasible solutions of the problem. Then, the metaheuristic approaches, based on tabu search and simulated annealing framework, were proposed for the production execution optimization.

On Discrete-Continuous Scheduling Problems with the Mean Flow Time Minimization

Joanna Jozefowska and Jan Weglarz
(Institute of Computing Science, Poznan University of Technology, Poznan, Poland)

The problem of scheduling n independent, nonpreemptable jobs on m parallel, identical machines is considered, when each jobs simultaneously requires a continuous, renewable resource. The processing time of job i, i = 1, ..., n, is a continuous, increasing and concave function of the amount of the continuous resource allotted to this job at time t. Each job is characterized by its processing demand and its processing rate function f. The processing model of a job defines the processing rate function as the first derivative of the state of job i at time t over time. The goal is to minimize the total completion time. The problem can be solved using a two-stage procedure. At the first stage a feasible sequence is created, while in the second step an optimal resource allocation is found solving, in general, a nonlinear mathematical programming problem.
The following results hold:
- an optimal continuous resource allocation does not depend on the processing demand of jobs,
- for a given feasible sequence of jobs, an optimal resource allocation can be found analytically for some classes of processing rate functions, e.g. power processing rate functions with the power 1/a, where a is from the set {1,2,3,4}, solving recursively n systems of at most (m+1) equations,
- an optimal feasible sequence is defined in the case when the processing rate functions of all jobs are identical,
- in the case with m=2 and f=sqrt(u), an optimal resource allocation can be found explicitly,
- it is an open question whether an optimal sequence can be found in polynomial time for arbitrary (concave) processing rate functions.

Job Due Date Setting in Manufacturing

Xiaoming Li and V "Sri "Sridharan
(Department of Management, Clemson University)

In this talk we present a new, general, parametric job due date setting procedure and analytically determine the optimum value of the parameter. We consider a M/M/1 system operating under the First-Come-First-Served priority rule and attempt to minimize variance of lateness, mean tardiness, and percent tardy jobs. We show that for the system considered, the mean tardiness and percent tardy jobs are always lower for a slack-based due date setting rule compared to the total work content (TWK) rule. We also discuss simulation results examining the performance of the proposed due date setting rule and the TWK rule under a range of values for shop utilization and different priority rules such as SPT, EDD, and MOD. These results show also that a slack based due date setting rule produces lower percent tardy and variance of lateness.

Single Machine Scheduling with Due Dates Depending on Processing Times

Valery Gordon
(Institute of Engineering Cybernetics, National Academy of Sciences of Belarus, Minsk, Belarus),
Jean-Marie Proth
(INRIA-Lorraine, Metz, France)

Recent results on scheduling with due date assignment are considered for the due dates which depend on processing times of jobs. The due date assignment rules can be either SLK rule (where due date is obtained by adding a positive slack to the job processing time) or TWK rule (when due date is obtained by multiplying processing time by a multiplier common to all jobs), or can combine both rules in PPW rule (Processing-Plus-Wait). The results on the single machine scheduling are analyzed and summarized for the various objective functions (maximum tardiness; total weighted earliness-tardiness; earliness costs with no tardy jobs). Most results can be found in the report of Leibniz-IMAG Laboratory "Scheduling with Due Dates (Annotated Bibliography of Complexity and Algorithms)" by G. Finke, V. Gordon and J.-M. Proth which is available from: www-leibniz.imag.fr/LesCahiers/. The work is supported in part by INTAS (Project INTAS 00-217).

Complexity Results for Shop Problems with Transportation delays

Peter Brucker
(Universitaet Osnabrueck, Fachbereich Mathematik/Informatik)
T.C. Edwin Cheng
(The Hong Kong Polytechnic University)
Sigrid Knust
(Universitaet Osnabrueck, Fachbereich Mathematik/Informatik)
Natalia V. Shaklevich
(University of Greenwich, School of Computing and Mathematical Sciences)

We consider shop problems with transportation delays where not only the jobs on the machines have to be scheduled, but also transportation of the jobs between the machines has to be taken into account.
Jobs consisting of a given number of operations have to be processed on machines in such a way that each machine processes at most one operation at a time and a job is not processed by more than one machine simultaneously. Transportation delays occur if a job changes from machine to another. The objective is to find a feasible schedule which minimizes some objective function.
A survey of known complexity results for flow-shop and open-shop environment is given and some new complexity results are derived.

Cylindrical Chromatic Scheduling of Dedicated 2-Processor Tasks

Marek Kubale and Adam Nadolski
(Technical University of Gdansk, Poland)

It occurs in many automated production systems that production proceeds in a periodic way. If the production cycle consists of several jobs of equal length, each requiring simultaneous use of two prescribed machines, then this situation can be modeled by a graph whose vertices correspond to machines, edges to jobs, and circular edge coloring to a cyclic schedule. It the talk we consider the complexity of two subproblems of this type: Pc|UET|Cmax and Oc||Cmax. We give proofs of NP-hardness of both subproblems and polynomial algorithms in some special cases restricting the number

of machines to m<4 and the structure of the associated scheduling graph. In particular, we show that both problems Oc3||Cmax and Occ3||Cmax are NP-hard.

Scheduling and Routing in Queuing Systems: Restless Bandit Index Policies
Jose Nino-Mora
(Department of Economics & Business, Universitat Pompeu Fabra, Spain)

In a seminal paper, Whittle (1988) proposed an appealing priority index policy for the hard restless bandit (RB) problem, which is only defined for some (indexable) models. Recently, we presented the first indexability conditions for restless bandits , based on the satisfaction by performance measures of partial conservation laws (PCL). We shall review the PCL-indexability conditions, their polyhedral foundation, our interpretation of them as a law of diminishing marginal returns, and several applications to the design of new index policies for routing and scheduling in queuing models with convex nondecreasing holding costs.

Airport Gate Management
Erwin Pesch and Ulrich Dorndorf
(FB 5 - Management Information Systems, University of Siegen, Siegen, Germany)

We present a novel modelling and solution approach for the task of assigning aircraft to terminal or ramp positions of an airport. The approach is used within a decision sup-port system that combines knowledge based and advanced mathematical optimisation techniques, thus offering the flexibility of a rule--based system while at the same time achieving a solution quality only obtainable with sophisticated optimisation algorithms. The optimisation model considers multiple objectives, including the maximisation of a total gate assignment preference score, the best use of towing opportunities with the minimum required number of towing operations, and the minimisation of deviations from a reference gate plan. A comparison of the results of the combined system with a traditional rule based approach shows dramatic improvements in solution quality.

Scheduling Manufacturing Systems via Fluid Model Optimization by Solving Continuous Linear Programs
Gideon Weiss
(Department of Statistics, The University of Haifa, Israel)

We present an algorithm for the exact solution of a separated continuous linear program in a finite number of simplex pivot steps, a problem which has been unsolved for ~50 years. The algorithm is based on the discovery that although the problem is to look for a vector of some general measurable functions, the actual solution has a representation which is finite, combinatorial structure. We discuss application of this to scheduling and control of manufacturing systems, e.g. semiconductor wafer fabrication, via fluid approximation.

Semi-Online Multiprocessor Scheduling with Given Total Processing Time
Hans Kellerer
(Institut fur Statistik und Operations Research, Universiteat Graz, Graz, Austria)
Vladimir Kotov

(Byelorussian State University, Faculty of Applied Mathematics and Computer Science, Minsk, Belarus)

We are given a set of identical machines and a sequence of jobs, the sum of whose weights is known in advance. The jobs are to be assigned on-line to one of the machines and the objective is to minimize the makespan. An algorithm with performance ratio 1.6 and a lower bound of 1.5 is presented. These results improve on the recent results by Azar and Regev, who proposed an algorithm with performance ratio 1.625 for the less general problem that the optimal makespan is known in advance.

Single Supplier Scheduling for Multiple Deliveries

Mikhail Y. Kovalyov
(National Academy of Sciences of Belarus, Minsk, Belarus)

In the problem under study, there is one supplier of intermediate components for several finished goods manufacturers. Whenever supplier's line switches from production of a component for one manufacturer to a component for another manufacturer, a sequence dependent setup cost occurs. Components are delivered to the manufacturers in batches of the same size. Processing and delivery times as well as start times for the manufacturers are given. The objective is to find a schedule such that the total setup cost is minimized, subject to the continuous production for all manufacturers. The problem is proved to be NP-hard. A dynamic programming algorithm for the general case and an algorithm linear in the total number of components for a special case of two manufacturers are developed. This research is supported by INTAS under grant number 00-217.

Designing a Timetable under Strict Legal Constraints

Carlos S. Azmat and Marino Widmer
(University of Fribourg, Fribourg, Switzerland)

In the current economical and industrial conditions, with the fluctuation of the demand, designing a timetable to define a work schedule for each employee is not an easy task. Moreover, due to legal and cost constraints, it is not always possible to engage and dismiss the employees according to the production requirement. In this paper, a three-step method to assign the daily work for full-time employees working on one shift is presented. It takes into account a set of strict legal constraints (in particular the one concerning the holidays) and it guarantees that the defined workforce is minimal, when it is assumed that the employees are able to perform each task in the workshop.

A Stochastic Scheduling Model for Software Projects

Frank Padberg
(Fakulteat fur Informatik, Universiteat Karlsruhe, Germany)

Staff is the most valuable resource today in software development. In view of the shortage of software developers, it is more important than ever that software project managers plan and schedule their development projects in such a way that the developers are deployed as effectively as possible. Planning and scheduling a software project is extremely difficult though:
- The time needed to complete a particular software development activity is known only roughly. The time needed depends on technical factors such as the complexity of the piece of code to develop, but also on human factors such as the experience of the developers.

- Software is an immaterial product. Thus, tracking the actual progress of a software project is difficult, making it hard for a manager to tell when it's time to take controlling action such as reassigning tasks.
- It is typical for software projects that activities which run in parallel interfere with each other. Feedback between activities often leads to unanticipated rework and delays.

We present a scheduling model which is tailored to the dynamics of software projects. In the model, activity times are stochastic, resources are constrained and non-identical, and scheduling is preemptive. As opposed to other stochastic project scheduling models, the completion time of an activity explicitly depends on the amount of feedback received from concurrent activities. As a consequence, the completion time of an activity also depends on the scheduling strategy.

Technically, the model is a discrete-time Markov decision process. The corresponding stochastic optimization problem of minimizing the expected makespan can be solved using stochastic dynamic programming. We outline how to exploit simulation in order to cut down the computing time of the optimization algorithm. We also examine the link between our optimization problem and simulation-based optimization techniques in machine learning.

Real-Time Node Allocation on Commodity-Based Supercomputers

Michael Bender (SUNY Stony Brook)
Cynthia Phillips (Sandia National Laboratory)
Vitus Leung (Sandia National Laboratories), Esther Arkin (SUNY Stony Brook), David Bunde (University of Illinois), Jeanette Johnston (Sandia National Laboratories), Alok Lal (Tufts University), Joseph S.B. Mitchell (SUNY Stony Brook), Steven Seiden (Lousiana State University)

The Computational Plant or Cplant is a family of commodity-based supercomputers under development at Sandia National Laboratories. This talk describes resource-allocation strategies to achieve processor locality in Cplant and other supercomputers. Users of Cplant and other Sandia supercomputers submit parallel jobs to a job queue. When a job is scheduled to run, it is assigned to a set of processors. To obtain maximum throughput, jobs should be allocated to localized clusters of processors to minimize communication costs and avoid bandwidth contention caused by overlapping jobs. This talk introduces new allocation strategies and performance metrics for general topologies. In particular, we order the processors so that processors that have similar ranks in the order are physically close in the network. We then solve a one-dimensional allocation problem using bin-packing algorithms and greedy minimization of span (the maximum difference between the ranks of any two processors allocated to a job). For 2D topologies we can find a processor ranking using space-filling curves. For general topologies, we can use an integer program. We present preliminary results for simulations and a Cplant implementation which indicate that both space-filling curves and one-dimensional packing can improve processor locality. These new allocation strategies are implemented in the new release of the Cplant System Software, Version 2.0, fully installed in May 2002. We also state hardness results for minimizing the stretch of maximum span.

Parallel Processor Scheduling with Limited Number of Preemptions

Oliver Braun and Günter Schmidt
(Department of Information and Technology Management, Saarbrücken, Germany)

We compare the makespans of preemptive and i-preemptive schedules for the problem where a set of n independent jobs (no ordering constraints between them) has to be scheduled on m identical processors that operate in parallel. The objective consists in minimizing the makespan, i.e., the maximum completion time of any job. We speak of an i-preemptive schedule, if the maximum number of

preemptions is bounded above by a nonnegative integer number i. A preemtive schedule is allowed to interrupt a job and later resume the execution of the job without any loss of processing time yet performed. Also, the minimum time slice for preempting a job may be arbitrarily small. In nonpreemptive schedules, once a job is started, it is executed to completion without any interruption. It is easy to construct an optimal preemptive schedule with McNaughton's Wrap Around rule, whereas it is an NP-hard problem to construct an optimal nonpreemptive schedule. We show that the ratio of the optimal i-preemptive makespan C_max^ip* versus the optimal preemptive makespan C_max^p* is bounded above by C_max^ip* <= (2-2/(m/(i+1))+1) C_max^ip*. Furthermore, we show that the ratio of the length C_max^LPT of a nonpreemptive schedule following the LPT rule versus the optimal preemptive schedule length C_max^p* is bounded above by exactly the same bound for i=0. We give an example which shows the tightness of the bounds.

Relaxing Multiprocessor Scheduling to Multibin Packing Problems
Pierre Lemaire, Gerd Finke and Nadia Brauner
(Laboratoire Leibniz-IMAG, Grenoble, France)

In classical bin-packing problems, objects of a given height are to be packed into bins, either to minimize the maximum height of the bins or to use as few bins (of given size) as possible. An extension of this concept is studied by considering multibin objects that fill several bins to a certain height at the same time. This is also a relaxation of the well-known area of multiprocessor scheduling: simply turn the Gantt chart by 90° and let the multiprocessor tasks drop in order to absorb all idle times. In this way, processors become bins and multiprocessor tasks become multibin objects. There are many applications of this model, for instance in the design of SONET rings in telecommunication, in expertizing documents, file backups and in testing medical treatments on volunteers with certain drug incompatibilities (allergies). For the analysis of multibin problems, one may use the same attributes, known from the theory of multiprocessor scheduling: size, any, fix and set. Complexity results and exact and approximation algorithms are presented. Relaxing multiprocessor solutions yield a performance factor of two when compared to the optimal multibin solutions. The problem Bm/sizej/Hmax are easily solved if all heights are equal to one and solved by a form of McNaughton wrap-around method. For general heights, one may consider Best-Fit procedures. Set-problems are already NP-hard if the cardinality of setj is equal to two. Any-problems can be solved in pseudo-polynomial time for any number m of bins by a dynamic programming approach, in contrast to the corresponding mutiprocessor scheduling problems. These problems are pseudo-polynomial for m = 2 and 3, the problem is open for m = 4 and strongly NP-hard for m = 5 machines.

Scheduling of Investment Projects in Manufacturing Systems
Jürgen Zimmermann (TU-Clausthal)
Christoph Schwindt (University of Karlsruhe)

Project managers are frequently confronted with the problem to decide on whether some investment project should be performed or to select one out of several mutually exclusive investment projects from a given portfolio. For the assessment of investments, the net present value criterion is well-established in research and practice. In classical investment theory, investments are specified by a stream of payments, i.e., a series of payments with associated payment times. For a stream of payments, the net present value is obtained by summing up all payments discounted to time 0.
In case of investment projects, the payment times are no longer given in advance but are subject to the scheduling process. An investment project consists of a set of events each of which is associated with a paying in or paying out. Moreover, there are prescribed time lags between the occurrence of the events.

Thus, the stream of payments results from maximizing the net present value of the project subject to the given time lags. The formulation of this optimization problem presupposes the knowledge of the proper interest rate for discounting the payments and the specification of a maximum project duration (project deadline). In practice, however, often neither the exact interest rate nor a project deadline are known. The proper interest rate is a theoretical value and can only be estimated. The project deadline is generally the result of negotiations between the investor and the contractor. The purpose of this paper is to develop a parametric optimization approach which provides the maximum project net present value as a function of the interest rate and project deadline chosen. The resulting net present value curve then serves as a basis for the decision of the investor, which depends on his individual risk preference.

Unit-Time Job-Shop Scheduling via Mixed Graph Coloring

Yuri N. Sotskov
(Institute of Engineering Cybernetics, Minsk, Belarus)
Alexandre Dolgui
(Laboratory of Industrial Systems Optimization, University of Technology of Troyes, Troyes, France)
Frank Werner
(Otto-von-Guericke-University, Faculty of Mathematics, Magdeburg, Germany)

A unit-time scheduling problem with makespan criterion may be interpreted as an optimal coloring of the vertices V of a mixed graph (V, A, E). A function is called an optimal coloring if the number of used colors is minimal, the start vertex of each arc in A has a smaller color than the end vertex, and the vertices connected by an edge in E are colored by different colors. We consider an optimal coloring of a mixed graph G which defines a schedule minimizing makespan for the unit-time job-shop problem. For such a mixed graph G = (V, A, E), it follows that the resulting digraph without edges is the union of disjoint paths and the resulting graph without arcs is the union of disjoint cliques. We present complexity results for the case of short paths and small cliques, and for the case of long paths or large cliques. For the case of two paths, we improve a geometrical algorithm. For the general case of a unit-time job-shop problem, we develop branch-and-bound algorithms and test them on randomly generated mixed graphs of order up to 200 for the exact solution and of order up to 900 for the approximate solution.

Accept or Reject? The Odds Are…

Han Hoogeven
(Utrecht University)

We consider the single-machine scheduling problem of minimizing the number of late jobs. We first review and reinterpret the famous algorithm by Moore and Hodgson for the case of deterministic processing times as a dynamic programming algorithm. We then look at four problem classes with stochastic processing times. The first one has processing times that consist of a deterministic part and a random component that is independently, identically distributed for each job. The jobs in the other three classes have processing times that follow: (i) A gamma distribution with parameters α_j and β, where β is common to all jobs; (ii) A negative binomial distribution with parameters s_j and $\mathbf{p}$, where $\mathbf{p}$ is the same for each job; (iii) A normal distribution with parameters ν_j and δ_j^2. In this scheduling environment, the completion times will be stochastic variables as well; under these circumstances, we qualify a job as being on time if the probability that it is completed by the deterministic due date is at least equal to a certain given minimum success probability. We show that for the first, second, and

third class of instances the problem can be solved in O(n\log n) time, where we need the additional assumption of equal minimum success probabilities in the first case. For the case with normally distributed processing times we present a pseudo-polynomial time algorithm, and we prove that this is the best we can hope for by establishing weak NP-hardness. We also show that the problem of minimizing the weighted number of late jobs can be solved by an extension of the dynamic programming algorithm in all four cases; this takes pseudo-polynomial time. We further indicate how the problem of maximizing the expected number of on time jobs (with respect to the standard definition) can be tackled if we add the constraint that the on time jobs are sequenced in order of nondecreasing due dates**.**

Scheduling and Routing in Adversarial Queuing Networks

Jay Sethuraman
(Columbia University, New York, USA)

Adversarial queuing networks serve as a convenient tool for modeling packet injections in modern communication networks. This model combines important aspects of two traditional ways of modeling input traffic: the stochastic model, and the online model. In this talk we construct simple discrete review policies to route and sequence packets so as to minimize the number of packets in the system. We show that the total number of packets in the system under our policy is O(1/(1-r)), where r (0 < r < 1) is the rate of the adversary. This improves an earlier result of Gamarnik (STOC 1999), which proves a bound of O(1/(1-r)^2).

Complexity Results in a Just in Time Scheduling Problem

Nadia Brauner
(Laboratoire Leibniz-IMAG, Grenoble, France)
Yves Crama
(Ecole d'Administration des Affaires, Université de Liege, Liege, Belgium)

This talk revisits the maximum deviation just-in-time (MDJIT) scheduling problem previously investigated by Steiner and Yeomans (1993).
An instance of this problem consists of a number n of different part types and of the integer demand di for part type i (i = 1, 2 ... n). All part types are produced on the same equipment (typically, a mixed-model assembly line) and the production of each part requires one unit of time. We denote by ri=di/sum(di) the ideal but fractional production rate for parts of type i. The term ``ideal" refers here to the fact that, at each instant, we would like the line to have assembled part type i in proportion ri. Such a schedule would be uniformly ``leveled". Obviously, perfectly leveled schedules are never attainable, but the aim of JIT control systems is to keep the actual production of each part as close as possible to its ``ideal rate". Monden (1983) states that this is a main objective of Toyota's JIT systems.
We recast this problem as a matching problem in a bipartite graph and we derive necessary and sufficient conditions for the existence of a schedule with a given objective function value. We use the previous result to establish that the MDJIT problem is in Co-NP and to prove that the problem can be solved in polynomial time when the number of part types is fixed. We then propose some conjectures concerning the structure of instances with small deviation.

Scheduling Divisible Tasks on Heterogeneous Processors
Pierre-Francois Dutot
(ID-IMAG, Grenoble, France)

The divisible load task model was introduced by Cheng and Robertazzi in 1988 to deal with highly parallel applications. The tasks can be splitted in any number of independent pieces of any size which can be scheduled concurrently on different processors. In the last 15 years the problem has been extensively studied for different kinds of processors topology. Recently some works on heterogeneous processors have been published as in IPDPS 2002. The main contribution of the talk is a polynomial algorithm to schedule n identical tasks on a chain of heterogeneous processors.

A New Clustering Algorithm for Large Communication Delays
Renaud Lepere and Denis Trystram
(ID-IMAG, Grenoble, France)

The problem of clustering is an alternative of usual scheduling algorithms. It consists of gathering the tasks of a parallel application with no limitation on the number of processors. It has been extensively studied as a basic step for finding efficient scheduling methods.
We present here a new method based on a recursive decomposition of the task graph. We first establish a general bound on a class of clustering with nice properties called convex clustering, then, we propose an algorithm for building convex clusters by finding two sets of independent tasks as large as possible that are possible to be executed in parallel and apply recursively the same decomposition on each set. Some preliminary experiments are presented for assessing the interest of this method.

Mixed Integer Programs for Scheduling Printed Wiring Board Assembly in Surface Mount Technology Lines
Tadeusz Sawik
(University of Mining and Metallurgy, Cracow, Poland)
An exact approach by mixed integer programming is proposed for scheduling printed wiring board assembly in SMT (Surface Mount Technology) lines. A SMT line consists of several processing stages in series, separated by finite intermediate buffers, where each stage has one or more identical parallel machines. The objective is to determine an assembly schedule for a mix of boards, so as to complete the boards in minimum time. The scheduling problems with machines continuously available and with limited machine availability are considered and, in addition, batch vs. cyclic scheduling modes are compared. The proposed approach can be used for optimization of assembly schedules in the electronics manufacturing by using commercially available software for mixed integer programming. This has been illustrated with a set of numerical examples modeled after real world SMT lines and some computational results with the CPLEX solver are reported .

Selfish Traffic Allocation
Artur Czumaj
(New Jersey Institute of Technology)
Berthold Voecking and Piotr Krysta
(MPI Saarbruecken, Germany)

We investigate the price of selfish routing in non-cooperative networks, like the Internet. In such networks, users may follow selfish strategies to optimize their own performance measure and therefore their behavior does not have to lead to optimal performance of the entire network. In this talk we investigate the worst-case coordination ratio, which is a game theoretic measure aiming to reflect the price of selfish routing.
Following a line of previous work, we focus first on the most basic networks consisting of parallel links with linear latency functions. Our main result is that the worst-case coordination ratio on m parallel links of possibly different speeds is exactly of order log(m) / logloglog(m).
Next, we present the first thorough study of this model for general, monotone families of cost functions and for cost functions from Queuing Theory. Among others, we give a precise characterization of cost functions having a bounded/unbounded coordination ratio, we show that an unbounded coordination ratio under the min-max objective implies an unbounded coordination ratio under the total-latency objective and vice versa.

A Buffer Minimization Problem for the Design of Embedded Systems
Alix Munier Kordon
(Université Pierre et Marie Curie)

We consider a set of n tasks, each of them is composed by a set of sequential operations. A set of buffers B is given: each buffer b from B is defined between two tasks Ti->Tj has a weight wb and is managed as a FIFO structure. A given set of operations from Ti write a data in the buffer b, and a set of operations from Tj get a data in b. The writings and readings on buffers generate precedence constraints between the operations. The limitation of the size of the buffers generates an other set of precedence constraints between the operations. Then circuits in this precedence graph may appear. The aim is to find the size of each buffer D(b) such that the sum of the terms D(b)wb on B is minimum and no circuit on the precedence graph occurs. We prove that this problem is polynomial for 2 tasks using a flow algorithm. We also prove that the problem is NP-hard in the strong sense for 3 tasks.

Incorporating Multitasking into a Static Scheduling Model
Frank Drews
(Technical University of Clausthal, Germany)

The problem of scheduling tasks in a multiprocessor system poses a fundamental problem in parallel processing. A large body of research efforts addressing these problems has been reported in the literature. We consider the objective of minimizing the completion time of a parallel application, represented by a directed acyclic graph (DAG), by properly allocating the tasks to the processors at compile time. Our goal is to introduce a basic scheduling model for multiprocessor systems with operating systems that run in multitasking mode and fulfill these conditions. We start to formalize the processing of tasks in multitasking operating systems and introduce a simple queuing model. In this model a (off-line) scheduling-strategy of an operating system scheduler can be represented formally as a function and describes how to assign time-slices to tasks. On the basis of these model we introduce the multitasking scheduling model (MT model). In this model a task is assigned to the queue of a processor and stays in this queue until finished. The execution of one ore more tasks can be overlapped. A function specifying the operating-systems scheduling strategy is a parameter of this model. We prove the NP-completeness of the general problem and present some bounds for optimal schedules. Furthermore we show how to represent MT schedules efficiently which is also a basis for algorithmic approaches. As the basic MT model does not incorporate the overlap of task processing

and I/O we extend the model by an additional I/O resource. For this problem it can be shown by reduction from 3-Partition that even problems with one processor and without precedences are NP-complete. Finally we present some experiments to compare our results with measured values from a concrete symmetric multiprocessor system.

A Branch-and-Bound Procedure to Minimize Total Tardiness on One Machine with Arbitrary Release Dates

Philippe Baptiste, Jacques Carlier, Antoine Jouglet
(Universite de Technologie de Compiegne, France)

In this paper, we present a Branch-and-Bound procedure to minimize total tardiness on one machine with arbitrary release dates. We consider the scheduling situation where n jobs J_1, …, J_n have to be processed by a single machine and where the objective is to minimize total tardiness. Associated with each job J_i, are a processing time p_i, a due date d_i, and a release date r_i. A job cannot start before its release date, preemption is not allowed, and only one job at a time can be scheduled on the machine. The tardiness of a job J_i is defined as $T_i = \max\{(0 , C_i - d_i)\}$, where C_i is the completion time of J_i. The problem is to find a feasible schedule with minimum total tardiness $\sum T_i$. The problem, denoted as $1 \mid r_i \mid \sum T_i$, is known to be NP-hard in the strong sense. We introduce new lower bounds and we generalize some well-known dominance properties. We also describe new dominance properties and propagation rules for this problem along with intelligent backtracking techniques. Our procedure handles instances as large as 500 jobs although some 60 jobs instances remain open. Computational results show that the proposed approach outperforms the best known procedures. Several techniques have been generalized to other criteria such as total weighted tardiness or total weighted completion time.