Jacek Blazewicz, Klaus Ecker(editors):

Scheduling in Computer & Manufacturing Systems

Dagstuhl-Seminar-Report; 114
15.05.-19.05.95 (9520)

ISSN 0940-1121

Tel.: +49 - 6871 - 2458
Fax: +49 - 6871 - 5942

Das Internationale Begegnungs- und Forschungszentrum für Informatik (IBFI) ist eine gemeinnützige GmbH. Sie veranstaltet regelmäßig wissenschaftliche Seminare, welche nach Antrag der Tagungsleiter und Begutachtung durch das wissenschaftliche Direktorium mit persönlich eingeladenen Gästen durchgeführt werden.

Bezugsadresse: Geschäftsstelle Schloss Dagstuhl
Universität des Saarlandes
Postfach 15 11 50
D-66041 Saarbrücken, Germany
Tel.: +49 -681 - 302 - 4396
Fax: +49 -681 - 302 - 4397
e-mail: office@dag.uni-sb.de
url: http://www.dag.uni-sb.de

INTERNATIONALES BEGEGNUNGS- UND FORSCHUNGSZENTRUM FÜR INFORMATIK

Schloß Dagstuhl

Seminar Report

Jacek Blazewicz, Klaus Ecker

Scheduling in Computer and Manufacturing Systems

May 15 - 19, 1995

Overview

During the week of May 15 - 19, 1995, the Seminar on Scheduling in Computer and Manufacturing Systems was organized by Jacek Blazewicz (Politechnika Poznanska, Poland) and Klaus Ecker (Technical University of Clausthal, Germany). Participants came from Universities or Research Centers from Austria (1), Belarus (2), Belgium (1), Canada (2), France (8), Germany (12), The Netherlands (3), Poland (7), Switzerland (1), and USA (4).

Altogether 39 lectures covered a large area of scheduling problems with various aspects of scheduling, both in algorithmic approaches and applications. Areas of particular interest were task scheduling under consideration of communication and shop problems.

The participants appreciated the outstanding local organization and the environment including all Dagstuhl facilities which not least enabled a successful workshop.

In the name of all participants,
Jacek Blazewicz, and Klaus Ecker

Reported by Martina Kratzsch

Participants

Jörg Apsel, TU Clausthal
Evripides Bampis, Université Paris-Sud
Jacek Blazewicz, Politechnika Poznanska
Jean Paul Boufflet, Université de Technologie de Compiegne
Pascal Bouvry, CWI Amsterdam
Peter Brucker, Univerität Osnabrück
Rainer Burkhard, TU Graz
Jean-Yves Colin, Université Paris VI
Yves Crama, Université de Liege
Moshe Dror, University of Arizona
Maciej Drozdowski, Politechnika Poznanska
Klaus Ecker, TU Clausthal
Gerd Finke, Université Joseph Fourier Grenoble
Lucian Finta, INRIA Sophia Antipolis
Valery Gordon, Belarus Academy of Science Minsk
Frederic Guinand, LMC-IMAG Grenoble
Jeet Gupta, Ball State University Muncie
Wolfgang Halang, FernUniversität Hagen
Claire Hanen, Université P. et M. Curie Paris
Reiner Hirschberg, PDV-Systeme Goslar
Adam Janiak, Politechnika Wroclawska
Jörg Kaiser, GMD St. Augustin
Roger Kieckhafer, University of Nebraska Lincoln
Mikhail Kovalyov, Belarus Academy of Science Minsk
Martina Kratzsch, TU Clausthal
Marek Kubale, Politechnika Gdanska
Wieslaw Kubiak, Menorial University of Newfoundland
Jan Karel Lenstra, Eindhoven University of Technology
Eric Luit, Eindhoven University of Technology
Rolf Möhring, TU Berlin
Jerzy Nawrocki, Politechnika Poznanska
Erwin Pesch, Universität Bonn
Jürgen Sauer, Universität Oldenburg
Günter Schmidt, Universität des Saarlandes, Saarbrücken
Philippe Solot , Ciba-Geigy AG, Basel
V. Sridharan, Clemson University
Eugeniusz Toczylowski, Politechnika Warszawska
Denis Trystram, LMC-IMAG Grenoble
Jan Weglarz, Politechnika Poznanska
Frank Werner, Otto-von-Guericke-Universität Magdeburg
Jia Xu, York University

Contents

JÖRG KAISER

DIRECT - An Adaptive Real-Time System: Concepts, Application, and Evaluation

ROGER KIECKHAFER

Using Automatically Generated GSPN's for Performability Modeling of Flexible Manufacturing Systems with Neuro-Fuzzy Controllers

MIKHAIL KOVALYOV

Parallel Machine Batch Scheduling: Algorithms, Complexity, and Approximation

MAREK KUBALE

Open Shop Problem with Zero-One Time Operations and Integer Release Time / Deadline Intervals

WIESLAW KUBIAK

Polynomial Time Algorithms for Unit-Time Job Shops

JAN KAREL LENSTRA

Lower Bounds for the Multiprocessor Flow Shop

JAN KAREL LENSTRA

Open Shop Scheduling: 3 is Easy, 4 is Hard

ERIC LUIT

Specification of Timing Requirements with Timed Precedence Constraints

ROLF MÖHRING

Bounds for the Project Duration Distribution in Stochastic Project Scheduling

JERZY NAWROCKI

Maximum Execution Time Prediction of Multi-Module Software

ERWIN PESCH

Grouping in Flexible Manufacturing

JÜRGEN SAUER

Knowledge-Based Scheduling and Meta-Scheduling

GÜNTER SCHMIDT

Office Scheduling

PHILIPPE SOLOT

Production Planning in the Chemical Industry

V. SRIDHARAN

Tactically Delayed Scheduling

EUGENIUSZ TOCZYLOWSKI

On Temporal Decomposition Algorithm for Scheduling Preemptible Tasks on Parallel Processors

DENIS TRYSTRAM

Minimizing the Overhead for Tree-Scheduling Problems

JAN WEGLARZ

On Discrete-Continuous Scheduling

FRANK WERNER

Heuristics for Shop Scheduling Problems with Batch Setup Times

JIA XU

Scheduling in Hard Real-Time Systems

Abstracts

Scheduling with Large Communication Delays

by EVRIPIDIS BAMPIS (joint work with Jean-Claude König)

We consider a task graph model where computational problem is represented by a directed acyclicgraph (dag). The basis of our analysis is the communication delay between the time some result has been produced on some processor and the time this information can be used by another processor. This communication delay, denoted by c is the fundamental parameter of a parallel architecture. Given this value, the problem that we have to solve is how to schedule the dag on the processors under the following communication model:

if task i starts its execution at time t on processor P_1 and task j is a predecessor of i in the dag, then either j starts its execution not later than time $t - p_j$ on processor P_1 or not later than time $t - p_j - c$ on some other processor, where p_j represents the execution time of task j.

Many papers studied the complexity of this problem in the case where for every task of the dag $p_j = 1$, and $c = 1$ ($P \mid prec, c = 1, p_j = 1 \mid C_{max}$). This problem was first adressed by Rayward-Smith who established the *NP*-hardness of the problem. Picouleau considered the same problem and showed that the problem of deciding if a dag can be scheduled on at most 3 units of time is polynomial. He had also considered the case where the number of processors is unbounded ($\overline{P} \mid prec, c = 1, p_j = 1 \mid C_{max}$) and he established the NP-completeness of the problem of deciding if an arbitrary dag can be scheduled in at most 8 time units. Independently Hoogeveen Lenstra and Veltman, and Saad showed that this last problem becomes NP-complete for $C_{max} = 4$. A corollary of this resultis that there will never exist a polynomial-time algorithm with performance bound smaller than $\frac{5}{4}$, unless $P = NP$. Similar results have been established for the case where the number of processors is unbounded and the performance bound becomes $\frac{7}{6}$. In this talk, we are interested in the case where the execution time of the tasks is equal to one and the communication delays a constant value greater than 1; does there exist a schedule of length L?

An Exact Method for Optimal Decomposition and Scheduling of an Application on a Bus Oriented Master-Slaves Multiprocessor System

by JEAN PAUL BOUFFLET (joint work with Jacques Carlier)

A bus oriented multiprocessor contains one bus system to which all the processors are interconnected. It is supposed that there is one master processor. The most serious problem is the bus bottleneck. Only one slave can establish communications through the bus with the master. A second problem is to maintain busy each processor. It is supposed we have to execute an application which uses a big amount of data and which can be split into several independent processes. Moreover the processing time is a linear function of its amount of data. The objective is to minimize the makespan by taking into account communication durations and processing times. At first the properties of optimal schedules are studied. Next we draw up a branch and bound method. The bound is obtained by relaxing some constraints. One dominance property is proved in order to prune the tree. The branching is based on the scheduling of sending subtasks and returning results. Computational results are reported showing the efficiency of the method. To conclude we expose the interest of this approach for finite elements decomposition methods.

Mapping Undirected Task Graphs on DMPC

by PASCAL BOUVRY

Efficient use of distributed-memory parallel systems requires the use of specific tools of which mapping is one of the most important ones. The aim of the mapping tools is to minimize the execution time of parallel programs on distributed-memory machines by controlling the use of computation and communication resources.

A parallel program can be represented as a graph of interconnected tasks which can be executed in parallel. The tasks execution times and inter-task communication costs of some regular programs can be entirely determined at compile time. In this case, it is possible to perform static task allocation in advance. This is known as the *mapping* operation whose complexity is exponential in the general case. This problem is known to be *NP*-Complete even in its most elementary form, the bin-packing problem (which consists simply of allocating independent tasks on a set of processors). Because of this complexity, numerous heuristic solutions have been proposed, representing different tradeoffs between computation cost and quality of mapping.

This talk presents a *mapping toolbox*, implementing several „classical" mapping algorithms. This toolbox was used to assess different cost functions by computing the relation between the value of these functions, optimized by the mapping algorithms, and the actual execution times of parallel programs. The implemented mapping algorithms were also evaluated, by comparing the execution times of a set of representative parallel programs mapped using the algorithms of the toolbox. The results of a large number of experiments were used to assess the cost functions by comparing their values to the actual execution times of the benchmarks. I will also talk about a new mapping method using monitoring tools to refine the predictions of tasks execution times and inter-task communication costs.

This work is part of the IMAGE-APACHE research project whose aim is the design and development of a general programming environment to balance automatically the load of parallel applications, resulting in reduced development time and increased portability of parallel applications.

Improving Local Search Heuristics for Some Scheduling Problems

by PETER BRUCKER (joint work with J. Hurink, F. Werner)

In many practical situations one is concerned with scheduling problems where a set of jobs has to be sequenced on some machines such that certain restrictions and conditions are satisfied and an optimal schedule with respect to some criterion function is to be determined. Since most of such scheduling problems are *NP*-hard exact algorithms such as branch and bound can solve these problems only if its dimension is small. Thus, heuristic algorithms are often applied to determine solutions that hopefully are not too far away from the global optimum.

Popular heuristics are local search methods. These methods are based on an underlying neighborhood structure. The choice of a suitable structure has some important influence on the quality of the search algorithm.

For a given scheduling problem with a finite set $\mathcal{S}$ of feasible solutions local search is a procedure that moves iterative through the set $\mathcal{S}$. In each step it goes from one solution $s \in \mathcal{S}$ to some 'adjacent' feasible solution. The possible moves from s to an adjacent solution are restricted by a set OP of possible operators op. For each $op \in OP$ the set $\mathcal{S}^{op}$ denotes a subset of $\mathcal{S}$ for which op is defined. Thus, $op \in OP$ is a function op: $\mathcal{S}^{op} \rightarrow \mathcal{S}$.

$$\mathcal{N}(s) := \{ \text{op}(s) \mid op \in OP, s \in \mathcal{S}^{op} \}$$

is the set of all possible **neighbors** of s. The sets

$$\mathcal{N}(s); \quad s \in \mathcal{S}$$

define a neighborhood on the set $\mathcal{S}$.

Whereas iterative improvement allows only moves from s to neighbors $s' \in \mathcal{N}(s)$ that have a better objective value, simulated annealing and tabu search allow moves to non-improving solutions s'. Especially, the last two methods have been applied successfully to many optimization problems. However, it often takes a large amount of computational time to get good results.

The objective of this work is to improve a given neighborhood $\mathcal{N}_1(s)$, $s \in \mathcal{S}$ by replacing the original set $\mathcal{S}_1$ of feasible solutions by a set $\mathcal{S}_2$ of solutions that are locally optimal with respect to a neighborhood $\mathcal{N}_1(s)$, $s \in \mathcal{S}_1$ (or even by a subset of all locally optimal solutions that contains at least one global optimum). The new neighborhood $\mathcal{N}_2(s)$, $s \in \mathcal{S}_2$ is constructed in such a way that is connected if $\mathcal{N}_1$ is connected. The constructions which are problem dependent are given for the following scheduling problems: $Pm \mid\mid C_{\max}$, $1 \mid\mid \Sigma T_i$, $1 \mid prec \mid \Sigma U_i$, and a class of scheduling problems with precedence constraints.

Numerical results are presented which document the usefulness of this method.

Minimization of Noise in Communication Networks

by RAINER E. BURKARD (joint work with E. Cela, T. Dudas and G. Wöginger)

We consider n communication centers which have to be assigned to the vertices of an undirected network. Each communication center transmits messages to each other center at a specified rate per time unit. In case that there is no direct connection between two centers the messages are routed via intermediate centers.

Once the communication centers are assigned to the nodes of the network he messages have to be routed. We distinguish between two different models for the routing pattern:

In the single path model for every pair of communication centers a single path is selected and all messages are routed along this path. In the fractional model the messages are split into several parts which may be routed along different paths.

The overall amount of traffic going through center i as intermediate center is called the noise in center i. Our goal is to find an assignment of the communication centers to the vertices of the network and a routing pattern that the maximum noise in a center becomes minimum.

We address complexity issues and point out efficiently solvable special cases for this problem. Moreover we show how lower bounds for the noise can be computed in an efficient way for use in a branch and bound method. We report the computational behavior of a branch and bound approach. Moreover neighborhoods are defined for use in local search algorithms. Simulated annealing and tabu search are compared with the results found by branch and bound.

Work in progress asks for a rational explanation for the shown behavior.

Scheduling Tasks and Communications on a Virtual Distributed System

by JEAN -YVES COLIN (joint work with P. Colin)

A Virtual Distributed System (VDS) is a system with an infinite number of homogenous processors, and a complete communication network between these processors.

In this talk, we present results on the problem of scheduling a task graph with communication delays on a VDU with limited capacity channels, that is each link between two processors contains two unidirectional channels, one in each direction, each able to carry one message at most.

We first present a polynomial algorithm that uses duplication and an extension of the CPM method to solve the problem if processing times are superior or equal to communication times. We then give some results when the length of the tasks and communications is not fixed but only bounded by min. and max. values. Finally, acknowledging the fact that there are no latest release dates in the usual sense in this kind of problem, we propose a notion of „no choice" latest release dates using some examples.

Cyclic Scheduling of Identical Parts in a Robotic Cell

by YVES CRAMA (joint work with Joris van de Klundert)

We consider a flow shop of m machines in which one type of product is to be produced. There are no buffers in the flowshop and the transportation of the items between the machines is taken care of by a robot. We consider the problem of finding the shortest cyclic schedule for the robot, that is a schedule that can be repeated infinitely many times and has maximum throughput rate. The problem has been solved by Sethi et al. (1992) in the case where $m \leq 3$. These authors also observed that, for general m, only $m!$ possible tours have to be considered for the robot. However, it is not at all trivial to find a shortest schedule given a tour for the robot. We show that the set of pyramidal tours, defined analogously to pyramidal tours for TSP, necessarily contains an optimal tour. Furthermore, we present an algorithm for finding the shortest cyclic schedule for any pyramidal tour. Finally, we show how these results can be used to solve the problem in polynomial time, using dynamic programming techniques. More specifically the recognition version of the problem can be solved in $O(m^2)$ time and its optimization version in $O(m^3)$ time.

„Strong-Weak“ Constrained Scheduling

by MOSHE DROR

We consider precedence relation in scheduling and distinguish between strong and weak precedence in chains and trees. We prove that for the case of chain precedence, the weak and strong precedence relations have nontrivial complexity implications. In some cases, a given scheduling problem with strong chain precedence is *NP*-hard while the weak chain case is easy, and for other problems the opposite is true. This observation can be extended to tree precedence relation (for in-trees and out-trees). Some scheduling problems with strong tree precedence relations are *NP*-hard while the weak tree problem is easy and there are problems for which the opposite holds. The above results indicate that when examining scheduling problems with precedence relations, we need to indicate more precisely the precedence structure the problem assumes since it might have important algorithmic implications.

Performance Bounds and Scheduling Divisible Tasks on Hypercube Multiprocessors

by MACIEJ DROZDOWSKI (joint work with Jacek Blazewicz)

In this work a problem of finding an optimal distribution of a divisible computational task among a set of processors is considered. In the model of parallel computer systems two important factors must be taken into account: speeds of processors and speeds of communication links. With regard to this, we propose a deterministic approach finding an optimal distribution of the job's load on a hypercube of processors. The method used allows also the determination of performance bounds on the hypercube architecture.

Single Processor Scheduling Problems with r-Width Bounded Task Precedences

by KLAUS H. ECKER

In this presentation a new method for optimally sequencing tasks on a single processor is introduced. Central idea is to construct an acyclic graph where each node represents a subset of pair-wise independent tasks of the given instance. This so-called *SIT*-graph is defined such that each path of maximum possible length represents a schedule, and vice versa. Depending on the scheduling objective, cost values are assigned to the edges of the *SIT*-graph. Finding an optimal schedule is then equivalent to determining a maximum length path of minimum total cost. In general, the time complexity of this method increases exponentially with the number of tasks. However, if the task precedence graph is of bounded width, the complexity depends polynomially on the number of tasks.

Hypergraph Approach for Scheduling Communication Demands in Networks

by KLAUS H. ECKER and REINER HIRSCHBERG

Usually in the past of scheduling theory it was assumed that only the processing of tasks consumes time. In fact looking on practical applications one has to realize that it is also important to take care of setup times, communication delays, transportation times, etc. We present a unique model based on hypergraphs to represent technical properties of networks with respect to communication. On the basis of the hypergraph representation of the network we develop several strategies for scheduling communication tasks. The same model can also be used for describing and solving automated guided vehicle problems.

General Flow Shop Models: From Job Overlapping to Job Deterioration

by GERD FINKE

In flow shop problems, the jobs consist of a sequence of operations. Usually the precedences are expressed in terms of operation completion times: the next operation can only start if the previous operation has been completed. We shall instead consider precedences in terms of starting times: only the start of subsequent operations is delayed. Thus we allow job overlapping which occurs in batch processing of production systems. A further application area concerns the sequencing of various projects. There the overlap tends the lengthen the processing time since only partial information from the previous stage is available. At the other extreme, whenever the waiting times between operations become too long, a job deterioration may occur, resulting also in longer processing times. This model has applications in the production of steel (where the material cools down with time), in emergency situations, and for maintenance and service operations. Detailed model descriptions are given and complexity results are reported.

Scheduling of Task Graphs with Communication Delays: Some Positive and Negative Results

by LUCIAN FINTA (joint work with Zhen Liu (all 3 problems), Ioannis Milis and Evripidis Bampis (last problem))

The first problem we consider is scheduling of parallel computations (task graphs) in multiprocessor systems with communication delays. Such a scheduling problem has recently been studied in the literature, mostly for the case where interprocessor communication times are fully determined (there are no contentions between messages). In this talk, we consider the scheduling problem with resource constraints. More specifically, the case where all interprocessor communications take place on a single bus is considered (only one message can be sent in any time slot). We show that even for very specific subproblems, the minimization of the makespan of task graph in such a single-bus multiprocessor system is *NP*-hard. Thus, the general scheduling problem of parallel programs with communication resource constraints is *NP*-hard. We consider several variants of the problem: tasks with preallocation or not, communications with independent-data semantics or common-data semantics. Our results are extended to the cases of blocking communications and of broadcasting communications, and can be applied to multiprocessor systems with shared memory.

The single-machine scheduling with precedence delays is the second problem we analyzed. A set of n tasks is to be scheduled on the machine in such a way that the makespan is minimized. The executions of the tasks are constrained by precedence delays, that is a task can start its execution only after any of its predecessors has completed and the delay between the two tasks has elapsed. In the case of unit execution times and integer lengths of delays, the problem is shown to be *NP*-hard in the strong sense. In the case of integer execution times and unit length of delays, the problem is polynomial, and an $O(n^2)$ optimal algorithm is provided. Both preemptive and non-preemptive cases are considered.

The scheduling of task graphs on two identical processors is the third problem considered. It is assumed that tasks have UET - UCT. The problem is to assign the tasks to the two processors and schedule their execution in order to minimize the makespan. A quadratic algorithm is proposed to compute an optimal schedule for a class of series-parallel graphs, which includes in particular in-forests and out-forests.

Batch Delivery Scheduling Problems

by VALERY GORDON (joint work with T. C. E. Cheng and M. Kovalyov)

Of obvious practical importance, batching problems became the subject of growing interest over last years. The paper considers batch delivery problems in a single machine shop to minimize the sum of batch delivery cost and total inventory holding cost. The problems are to determine simultaneously an optimal job partitioning into batches, optimal number of batches and job sequencing in each batch. Jobs in the same batch are to be delivered together upon the completion time of the last job in the batch. Jobs finished before this time have to wait until delivery and there are inventory holding costs for the jobs finished before the delivery date.

A polynomial algorithm for a special case of the problem is presented. The dynamic programming approach for minimizing total cost is considered.

The connection between the batch delivery problems and the parallel machine problems is established that gives the opportunity to use complexity results and algorithms obtained for parallel machine scheduling problems to study the considered problems.

Scheduling Intrees with Communication Delays on Identical Processors

by FREDERIC GUINAND (joint work with Christophe Rapine and Denis Trystram)

Consider that a program can be modeled by a precedence graph in which each task represents some computational parts and each arc a communication of data between two tasks. Within an execution model such that neither preemption nor duplication of tasks is allowed, but allowing the overlap of communication by computations, we propose an algorithm that yields optimal schedules for intrees with unit execution time tasks and unit communication time on two processors. The result is not new (Lawler, Veldhorst and Picouleau have the same result), but the strategy used is original. Moreover, this strategy (clustering) seems to be suitable for the same problem with other assumptions on the execution times or on the communication times. Indeed, the distance between the schedule provided by our algorithm and the optimal schedule for the problem $P2 \mid intree, p_i = \{1, 2, \ldots, k\}, c = 1 \mid C_{\max}$ is no more than $\frac{k}{2}$, and for the problem $P2 \mid intree, p_i = \{1, 2, \ldots, k\}, c \mid C_{\max}$, this distance is less than or equal to $\frac{c}{2}$.

For the more general case of the same problems but for a number of processors equal to m, we have shown that the bound given by Lawler $(m - 2)$ can be improved and that $\frac{(m-1)(m-2)}{2m}$ that constitutes a tight bound.

This work is a joint work with Christophe Rapine and Denis Trystram. We have produced moreover another algorithm for the m machines problem that has the same performance than Lawler's one. This algorithm is still based on a clustering strategy and seems to be suitable for intrees with execution times tasks greater than one and for intrees with communication delays constant greater than one.

Single Facility Scheduling with Two Operations per Job and Time-Lags

by JATINDER N. D.GUPTA

This paper considers the single facility scheduling problem where each job requires two operations separated by a minimum specified time-lag. The problem is shown to be NP-hard in the strong sense based on a transformation of the problem to a two-stage flowshop problem with time-lags. Approximate solution algorithms are developed to solve the problem with a worst-case performance ratio of $\frac{3}{2}$. Empirical evaluation shows that the performance of these algorithms is much better than the theoretical worst-case bound. Some special cases of this problem are also identified and solved.

Contemporary Research on Real-Time Scheduling Considered Obsolete

by WOLFGANG A. HALANG

The reason for research in scheduling is an economical one, *videlicet*, to optimize the utilization of resources. Up to the present time, almost all interest is directed towards processor scheduling. Departing from the requirements holding for real-time computing, in this paper it is shown that maximum processor utilization has become obsolete as an optimization criterion for industrial real-time systems. It is also shown that the earliest-deadline-first discipline and certain modifications thereof provide a satisfactory and final answer to all real-life scheduling needs. To this end, all intrinsic properties of this discipline are compiled and discussed in order to show that it is the most advantageous scheme at hand, characterized by efficiency and allowing for predictable system behavior. It is then pointed out how the method naturally extends to the scheduling of tasks having non-preemptable regions due to resource-access constraints. A sufficient condition is presented, which allows, at any arbitrary point in time and under observation of resource constraints, to check the feasible schedulability of the tasks competing for processor allocation. This condition applies to entirely non-preemptable tasks as well. Then, by taking industrial practice and actual cost relations into account, evaluation criteria and design principles for real-time computing systems are developed. The paper closes with pointing to those open optimization questions, scheduling research ought to address if it wants to deal with practically relevant problems, *videlicet*, minimization of software costs, software complexity, and complexity of schedules, synchronization sequences, inter-task communication, etc. In other words, simplicity is to be maximized to enhance system dependability and predictability of system behavior.

An Approximation Algorithm for Scheduling Dependent Tasks on Parallel Processors with Small Communication Delays

by CLAIRE HANEN

With the recent development of parallel architectures arised a new class of scheduling problems in which communication delays between different processors are considered. A precedence relation between two tasks represents a data transfer, the duration of which may depend on the mapping of the tasks over the processors.

This talk addresses the problem of scheduling a set of dependent tasks on parallel processors, assuming small communication delays (less than processing times). Although the problem is known to be *NP*-hard, even on an unlimited number of processors, optimal schedules have properties that can be fruitfully used in an approximation algorithm.

We first show that the notion of favorite successors, which has already been used for scheduling trees, induce a linear program with binary variables that models the problem assuming infinitely many processors. The continuous relaxation of this program is used to provide a solution with performance ratio bounded by $\frac{4}{3}$.

Then a new mechanism, taking into account given favorite successors, is included in a classical list scheduling scheme in order to build a schedule on m processors. The original feature of this mechanism is that it allows some ready tasks to be delayed, eventually producing idleness on a processor, in order to wait for a more important task.

The worst case analysis of this heuristic shows that its relative performance is bounded by $\frac{7}{3} - \frac{4}{3}m$ if the favorite successors are derived from the linear programming heuristic.

Scheduling Problems with Resources - Mathematical Models, Computational Complexity, Algorithms and Applications

by ADAM JANIAK

A short survey of scheduling problems with continuously – divisible resources (such as: computer memory, energy, gas, fuel, catalyzer, raw materials, financial outlay), i.e. with various (e.g. static, dynamic, etc.) models of job processing times and release dates, is given. Some practical examples of such problems are also given, e.g.:

- process in Soaking Pit and Blooming Mill Departments of Steel Mills with:
  - (i) minimization of the blooming mill standstills s.t. a given constraints on the gas flow intensity,
  - (ii) minimization of the total gas consumption s.t. a given level of production,
  - (iii) bicriterion approach - searching compromise points between blooming mill standstills and the total gas consumption;
- process in Forge Shop Department of Steel Plants,
- cooper electrorefining process in cooper industry,
- allocation of common virtual memory in multiprocessor system.

Next some examples of scheduling problem with resources are examined.

The first one is a single machine problem under the assumption that each job is available for processing at the moment (release date) which is some positive decreasing function with respect to the amount of continuously – divisible resource. This problem is considered for the following criteria:

- minimization of machine standstills s.t. a given constraint on the total amount of resource,
- minimization of the total resource consumption s.t. a given constraint on the maximum job completion time,
- bicriterion approach.

It is shown that the decision versions of the problems dealt with are strongly *NP*-complete. Some properties of the problems are presented. Polynomially solvable subcases of the problems and approximate algorithms (with worst case analysis) are also given.

The second problem is a permutation flow-shop one with the makespan criterion. It is assumed that the processing times of jobs on some machines are convex decreasing functions with respect to the amount of continuously – divisible locally and globally constrained resource. It is shown that two-machine problem is *NP*-hard even for identical slopes of linear models of operations on one of the machines and constant processing times on the other one. Due to the proved elimination properties the search tree of the proposed branch and bound algorithm is strongly restricted. Some approximate algorithms are also given.

DIRECT - An Adaptive Real-Time System - Concepts, Application, and Evaluation

by JÖRG KAISER (joint work with Hermann Streich, Martin Gergeleit)

DIRECT stands for Distributed Real-Time Control and incorporates concepts and mechanisms to construct real-time applications which are able to flexibly adapt to a changing environment. The talk concentrates on our work on dynamic scheduling which is prerequisite for any adaptation.

TPS (Task Pair Scheduling) is a dynamic scheduling mechanism which combines the idea of on-line guarantees with the notion of exception handling. On-line guarantees assure that a task once started has enough resources to successfully finish before a given deadline. This guarantee is given on the basis of Worst-Case-Execution-Times (WCET) of the respective task. TPS allows to use more optimistic assumptions like average (ACET) or optimistic (OCET) execution time. Obviously, a task with a given resource budget of ACET or OCET may not be able to finish all computations within this time. TPS deals with this situation by splitting a task into a so called main task and an exception task which is started when the resource budget of the main task is consumed without producing the expected results. A task pair is only scheduled if the execution of the exception task can be guaranteed and hence, the task pair can be completed in an application dependent well defined state.

As an example, a program for the remote control of a robot arm is presented. The TPS solution is evaluated using our monitoring facility which allows us to assess and visualize the fine-grained timing behavior of the application.

Finally, the concept of an adaptive real-time system is outlined, based on TPS and continuous system observation. Continuos system monitoring is used for on-line assessment of the system state to obtain the fine-grained timing information which allows us to determine the progress of an application. This information is then used to modify scheduling decisions.

Using Automatically Generated GSPN's for Performability Modeling of Flexible Manufacturing Systems with Neuro-Fuzzy Controllers

by ROGER M. KIECKHAFER (joint work with F. F. Choobineh, and Yi Xue)

A formal methodology is being developed for the specification and design of effective Neuro-Fuzzy controllers for Flexible Manufacturing Systems (FMSs). The goal is to maximize the expected values of long-term reward (performability) functions. System dynamics are modeled by a Generalized Stochastic Petrí Net (GSPN), hierarchically partitioned into a system subnet and a controller subnet.

Due to the size and complexity of real world FMSs, GSPN partitioning has been formalized to permit automated construction of the GSPN system subnet. A program has been written to generate a canonical GSPN for the system subnet given a simple description of the physical resources and the process plans of the jobs to be performed. This presentation describes the GSPN partitioning and construction process and the input language, and discusses the advantages of the approach.

Parallel Machine Batch Scheduling: Algorithms, Complexity, and Approximation

by MIKHAIL Y.KOVALYOV

A single machine batch scheduling problem can be formulated as follows. There are several independent non-preemptive jobs to be scheduled for processing on a machine. Jobs may be combined to form batches. A batch is a set of jobs processed jointly. A set-up time is required immediately before a batch is processed. A schedule is characterized by a partition of the set of jobs into batches, a sequence of batches and a sequence of jobs in each batch. The objective is to find a schedule minimizing a regular function of the job completion times. Two models for job completion times are considered. In the first, all jobs in a batch are completed together when the last job in the batch has finished processing. In the second, a job is completed immediately when its processing is finished.

For the case of parallel machines, besides the batching and sequencing decisions we have to find an allocation of the batches to the machines. The computational complexity of the above problem is studied. The *NP*-hardness proofs or polynomial time algorithms are presented for various special cases. For *NP*-hard problems, dynamic programming formulations are developed as well as fully polynomial approximation schemes. Open problems are indicated. Generalizations and some trends for future research are discussed.

Open Shop Problem with Zero-One Time Operations and Integer Release Time / Deadline Intervals

by MAREK KUBALE

The open shop scheduling problem has received considerable attention in the literature. Roughly speaking, the problem is to schedule n jobs consisting of different operations on m processors. Each of these processors performs a different type of operations. The operations of a job can be processed in any order but no two of them can be executed simultaneously. The aim is to construct a non-preemptive schedule with minimum makespan. The problem was extensively studied by Gonzalez and Sahni in 1976. They showed that finding an optimal schedule for this problem in *NP*-hard, if $m \geq 3$.

In this talk we consider a model of scheduling in which each execution time of an operation is either 0 or 1. This model of open shop remains still *NP*-hard, if a schedule with minimum total completion time is required. Therefore we assume the maximum completion time as a criterion of optimality. On the other hand, as in classical scheduling problems, we allow release times and deadlines for all jobs. Unfortunately, so specified open shop problem remains *NP*-hard. For this reason we assume that all release times and deadlines occur at integral times only. Due to such constraints the open shop problem considered here is easier to solve in extremal cases with respect to the number of unit time operations (UET), though it still remains *NP*-hard when there are at most three operations per job. We show that these extremal cases include the following:

(1) all possible mn operations have UET,

(2) at most $m + n$ operations have UET.

Next, the second algorithm is generalized to bounded cyclomaticity graphs. All the algorithms are capable of minimizing not only makespan but maximum lateness and maximum tardiness as well.

Polynomial Time Algorithms for Unit Time Job Shops

by WIESLAW KUBIAK

We study the problem of minimizing total completion time as well as the problem of minimizing makespan in two machine job shop with unit time operations. We propose efficient algorithms for the problems. The algorithms are polynomial with respect to a succinct encoding of the problem instances, where the number of bits necessary to encode a job with k operations is $O(\ log(k + 1)\ )$.

Lower Bound for the Multiprocessor Flow Shop

by JAN KAREL LENSTRA (joint work with Ann Vandevelde, Han Hoogeveen, and Cor Hurkens)

The multiprocessor flow shop is a flow shop in which each stage consists of a set of identical parallel machines instead of a single machine. As finding a minimum-length schedule is *NP*-hard, we set out to find good lower and upper bounds. Our main lower bound is based on the relaxation of the capacity of all machine sets except one. The resulting parallel machine problem is still hard, but further relaxation's give good bounds that can be calculated fast. Our upper bounds are, so far, obtained by relatively simple constructive rules. We investigate the theoretical relations between the bounds as well as their experimental performance.

Open Shop Scheduling: 3 is Easy, 4 is Hard

by JAN KAREL LENSTRA

In this tutorial lecture I presented some recent work done in cooperation with David Williamson, Leslie Hall, Han Hoogeveen, Cor Hurkens, Sergey Sevastjanov, and David Shmoys. My main purpose was to prove a negative result on the approximability of the optimum in open shop scheduling:

(1) Unless $P = NP$, there does not exist a polynomial-time algorithm which, for each instance of the open shop scheduling problem, finds a schedule that is guaranteed to be strictly shorter than $\frac{5}{4}$ time an optimal schedule.

This is a corollary to the following theorem:

(2) It is *NP*-complete to decide if an open shop has a schedule of length at most 4.

I also showed that our technique cannot be used to strengthen the negative result:

(3) It can be decided in polynomial time if an open shop has a schedule of length at most 3.

Theorem (2) was proved by a reduction from MONOTONE NOT-ALL-EQUAL 3-SATISFIABILITY. The proof of Theorem (3) is based on a reformulation of the problem as a restricted bipartite edge-coloring problem, which reduces to a matching problem.

Specification of Timing Requirements with Timed Precedence Constraints

by ERIC J.LUIT

Timing requirements of periodic control applications are more naturally expressed by Timed Precedence Constraints (TPC) than by periods, release dates and deadlines. A TPC is a precedence constraint $t_1 \xrightarrow{[s,l]} t_2$ between two tasks t_1 and t_2.This constraint requires that the start time $st(t_2)$ of t_2 satisfies $st(t_1) + s \leq st(t_2) \leq st(t_1) + 1$.

In a high-volume copier, the flash of an original and the release of a blank sheet are tightly related in time. This is not easily described by release dates and deadlines with respect to the beginning of a period. In many other applications, it is required that the response to an aperiodic event is completed before a deadline d. When static scheduling is used as in the Dependable Distributed Operating System (DEDOS), such response tasks are usually transformed to periodic activation's. The period p must then be chosen between $\frac{1}{2}d$ and $d - c$, where c is the worst-case execution time of the task. The length of the period depends on additional constraints imposed on an activation within the period. It is better to specify subsequent activation's by a TPC $\xrightarrow{[s,d-c]}$ with $c \leq s \leq d - c$. This leads to a lower average utilization and more flexibility for the scheduler.

The specification with TPCs matches closely with timing specifications that can be expressed in the DEDOS Application Language (DEAL). In DEAL, timing is specified by *time measurements* and *time requirements*. A time measurement defines the moment at which a statement is executed in the schedule. A time requirement specifies the time at which a statement must be executed relative to a time measurement. A pair (time measurement, time requirement) directly corresponds to a TPC.

Bounds for the Project Duration Distribution in Stochastic Project Networks

by ROLF H. MÖHRING (joint work with Rudolf Müller)

We present new results for evaluating a project network (= acyclic directed graph) $\mathcal{N}$ whose activities (= edges of $\mathcal{N}$) have random, possibly dependent, processing times. The aim is to obtain information about the random shortest project duration PD (= longest path length of $\mathcal{N}$). It is known (Hagstroem 88) that calculating the distribution function of PD at one point is NP-complete. This motivates - besides simulation - the construction of stochastic upper and lower bounds for the distribution of PD. We discuss existing approaches and elaborate on their relationship with combinatorial optimization (e.g. modifying $\mathcal{N}$ into series-parallel networks, time-cost tradeoff problems, k-longest paths). In theoretical respect, we unify and generalize several known bounds for independent processing times (e.g. Kleindorfer 71, Spelde 76, Dodin 85) by a new minor construction for networks and their path clutters. We also show that the evaluation problem at one point remains weakly *NP*-hard even for series-parallel networks with discrete processing time distributions, while it becomes polynomial in the maximum number of values of the project duration of a subnetwork of $\mathcal{N}$. We report on computational experience with these algorithms.

Maximum Execution Time Prediction of Multi-Module Software

by JERZY R. NAWROCKI

In hard real-time systems satisfying timing constraints is as important as satisfying functional specification. To guarantee that each task will meet its deadline, pre-run-time scheduling must be applied, which requires maximum execution time of each task to be known at compile time. Thus, the problem of maximum execution time prediction arises. Unfortunately, the solutions one can find in literature are not satisfactory. They either depend on running the program and measuring its execution time (that is not a technique suitable for time critical applications) or they are based on analyzing the program, but then recursion and dynamic data structures are disallowed. To overcome these difficulties it is proposed to split a hard real-time module into the following parts: interface, implementation, performance interface, and implementation description. The first two are well known. Performance interface describes the performance attributes that are available and implementation description contains a set of subroutines that compute the values of performance attributes for a given implementation.

Grouping in Flexible Manufacturing

by ERWIN PESCH

The traditional form of flexible manufacturing, i.e. the job shop, based on a functional layout of the production units, has been undergone rapid changes within the last two decades. Increasing competition, environmental conditions, increasing consumer demands, decreasing product life cycles, reduced delivery periods up to just-in-time etc. led to manufacturing technologies (FMS and group technology) in order to meet, basically, the objectives "increasing flexibility" and "decreasing fixed costs" (flow time, WIP etc.). Part type selection (choose a subset of the part types for immediate and simultaneous production), resource (machine or tool) grouping, production ratio determination (with respect to the selected part types), resource allocation and loading subject to technological and capacity constraints are the induced planning problems. We briefly review most popular approaches in order to divide the system into machine cells by grouping machines and parts. A new approach based on the similarity of part manufacturing characteristics and list appropriate encoding as well as some measure of machine similarity will be introduced. It leads to the problem of simultaneously solving clique partitioning like problems.

Knowledge-Based Scheduling and Meta-Scheduling

by JÜRGEN SAUER

Scheduling is one of the crucial tasks to be done in order to guarantee the efficient execution of the production processes. To support the human schedulers in performing this task two approaches are presented. The first one describes the requirements and a common architecture for knowledge-based scheduling systems and shows an example of an implemented system. Important features of a knowledge-based scheduling system are the integration of a graphical user-interface presenting data of the production area and supporting manual scheduling together with a heuristic scheduling algorithm incorporating the scheduling behavior and knowledge of the expert.

Inspired by the experiences with the knowledge-based scheduling system the second part describes a general approach for meta-scheduling which tackles the problems of creating and selecting the appropriate scheduling strategy in a given situation. This combines the use of meta-scheduling knowledge to select the appropriate strategy and the use of dynamic scheduling knowledge for the flexible creation of scheduling algorithms. A language is presented that allows the representation of scheduling and meta-scheduling knowledge. The implementation leads to a hybrid scheduling system, providing the description and integration of different scheduling approaches and the support in selecting the appropriate ones.

Office Scheduling

by GÜNTER SCHMIDT (joint work with Jörg Winckler)

Scheduling theory can be applied to various business problems where in the past manufacturing and logistics were the most popular. With office scheduling a new model for organizations will be presented in order to gain more efficiency in administration processes.

We will give a survey on selected scheduling problems in the office area and compare them to well known problems from classical scheduling theory. The main differences come from the observation that in office environments information has to be scheduled which can be shared by different processors utilizing arbitrary duplication or copy possibilities.

Taking this difference into account we will show that office problems in many cases are easier to solve and also the performance quality is better in comparison to classical problems from scheduling theory.

Production Planning in the Chemical Industry

by PHILIPPE SOLOT

In all branches of industry, the continuous increase in economical pressure and competition leads companies to search ways of improving the efficiency of their production. From a logistic point of view, this goal can be achieved by better organizing the material flow through the factories so as to optimize the utilization of the equipment and of the additional resources. Although the chemical and pharmaceutical industry mainly faces the same kind of production planning issues as the discrete manufacturing industries, particularities of the production processes to be considered tend to make the scheduling and planning problems more difficult.

This talk is a short presentation of two practical planning problems that we currently address as operations research consultants in a chemical-pharmaceutical company. We first briefly explain which are the typical characteristics of a chemical production environment. We then consider the issue of detailed scheduling within a factory and present a graphical interactive software package called *profiPLAN* that we developed to "solve" such problems. Finally, we present the pipeline planning problem which is an example of multi-level production planning in the pharmaceutical sector. We conclude the talk by some practice-oriented remarks concerning the problems presented.

Tactically Delayed Scheduling

by V. SRIDHARAN (joint work with Z. Zhou and J. J. Kanet)

Much of the existing job-shop scheduling literature concerns itself almost exclusively with non-delay (ND)schedules. ND schedules may not be appropriate for a number of cases. We present a taxonomy of scheduling problems and identify conditions when tactically delayed (TD) schedules may be a propitiate. TD schedules are defined as those which contain deliberate idle times between jobs. TD procedures are defined as those which permit TD schedules when appropriate. After reviewing the extant literature dealing with the so called E/T problem, we reach the conclusion that the approach taken to solve these problems - the approach of first finding the best sequencing and then inserting idle times - is inadequate and cannot guarantee optimality. Recognizing the NT-completeness of the problem, we propose a new heuristic approach called the DT approach which is based on decision theory. The DT approach is a dispatch like approach which takes a global view and is criterion independent. We present DT procedures for dynamic one-machine problems with regular and non-regular penalty functions. They permit TD schedules when appropriate. Experimental investigation indicates that the DT-TD procedures outperform several existing heuristics and find the optimum solution for a large fraction of problems studied. We believe that the DT approach is an attractive alternative whose effectiveness can be further improved by obtaining better estimates of job completion times.

On Temporal Decomposition Algorithm for Scheduling Preemptible Tasks on Parallel Processors

by EUGENIUSZ TOCZYLOWSKI

A two-phase aggregate / disaggregate algorithm for preemptive scheduling of independent tasks on general parallel processors with additional doubly constrained resources is presented. In this class of problems, one has to find a schedule minimizes a general, cost-based objective function, subject to conditions that various types of renewable and consumable resource limitations are satisfied by feasible schedules. The general limitations on resources are modeled by concave piecewise linear cost functions. A two-phase algorithm for solving this class of scheduling problems is based on aggregation and disaggregation of the aggregate solution by a temporal decomposition algorithm which generates a sequence of temporal detailed scheduling and resource allocation subproblems.

Minimizing the Overhead for Tree-Scheduling Problems

by DENIS TRYSTRAM (joint work with Evripidis Bampis, Frederic Guinand)

This talk is devoted to the study of the problem of scheduling tree precedence task graphs within the execution model introduced by Anderson-Beame-Ruzzo. It consists in minimizing the "overhead", defined as a trade-off between idle times and communication times. After a brief presentation of the most important complexity models, we examine the problem of scheduling binary trees on two processors and show that it is *NP*-hard. Then, we present an optimal algorithm for complete k-ary trees (for any k) on a given number of processors m.

On Discrete-Continuous Scheduling

by JAN WEGLARZ (joint work with Joanna Jozefowska)

Discrete-continuous scheduling problems arise when jobs simultaneously require for their processing at a time discrete and continuous resources. The simplest subclass of such problems contains the problems in which each job requires for its processing a machine from a set of m identical, parallel machines (a discrete resource) and an amount (unknown in advance) of a continuous, renewable resource (e. g. power). We consider job models in the form of a continuous, non-decreasing functions relating job processing rates to the resource amount allotted to this job at a time. Jobs are non-preemptive, available at the start of the process, and schedule length is the optimality criterion. We identify efficiently solvable subclasses, propose heuristics for the general case, and study their worst-case behavior.

Heuristics of Shop Problems with Batch Setup Times

by FRANK WERNER (joint work with Y. Sotskov, Th. Tautenhahn)

In the talk we consider shop problems with batch setup times. There are n jobs that are divided into G groups. We have to form batches consisting of jobs of the same group and to determine the optimal batch sequence on all machines. A setup occurs only before a new batch starts on a certain machine. In addition we consider different types of job availability, namely item and batch availability of the jobs. Contrary to the classical polynomial solvable shop problems with 2 machines and makespan minimization, the corresponding problems with batch setup times are already *NP*-hard. We consider different constructive and iterative heuristics for the permutation flow shop and the job shop problem with batch setup times. Most of the constructive algorithms are based on insertion techniques combined with some refinements (different beam search variants, parallel insertion of operation). Concerning iterative algorithms, we mainly consider the question of designing suitable neighborhoods for the case of problems with batch setup times but different meta-heuristics are also included. Results of a computational study for such problems are presented.

Scheduling in Hard Real-Time Systems

by JIA XU (joint work with David Lorge Parnas)

We are interested in large, hard-real-time systems with 50 - 500 processes with hard-real-time deadlines where the processes have different timing characteristics and interdependencies; high processor utilization is required; and no surprise timing failures are allowed. Examples of such systems include aircraft flight programs, nuclear plant control, traffic control, etc. There are two distinct approaches to scheduling processing hard-real-time systems. One is run-time scheduling, the other is pre-run-time scheduling. We explain why we believe that pre-run-time scheduling is essential if one wishes to guarantee that timing constraints will be satisfied in a large complex hard-real-time system. We examine some of the major concerns in pre-run-time scheduling and consider what formulations of mathematical scheduling problems can be used to address those concerns. We discuss further work that needs to be done to automate pre-run-time-scheduling.

E-Mail Addresses

Jörg Apsel	apsel@zaphod.in.tu-clausthal.de
Evripides Bampis	bampis@lri.fr
Jacek Blazewicz	blazewic@pozn1v.tup.edu.pl
Jean Paul Boufflet	boufflet@hds.uni-compieqne.fr
Pascal Bouvry	bouvry@cwi.nl
Peter Brucker	pbrucker@dosuni1.rz.uni-osnabrueck.de
Rainer Burkhard	rburkard@ftug.dnet.tu-graz.ac.at
Jean-Yves Colin	colin@lachesis.ipb.fr
Yves Crama	crama@vm1.ulg.ac.be
Moshe Dror	mdror@bpa.arizona.edu
Maciej Drozdowski	maciej_d@pozn1v.tup.edu.pl
Klaus Ecker	ecker@zaphod.in.tu-clausthal.de
Gerd Finke	gerd.finke@imag.fr
Lucian Finta	lucian.finta@sophia.inria.fr
Valery Gordon	mahaniok%bas10.basnet.minsk.by@demos.su
Frederic Guinand	guinand@imag.fr
Jeet Gupta	00jngupta@bsuvc.bsu.edu
Wolfgang Halang	wolfgang.halang@fernuni-hagen.de
Claire Hanen	hanen@litp.ibp.fr
Reiner Hirschberg	hirsch@zaphod.in.tu-clausthal.de
Adam Janiak	janiak@ict.pwr.wroc.pl
Jörg Kaiser	kaiser@gmd.de
Roger Kieckhafer	rogerk@cse.unl.edu
Mikhail Kovalyov	mahaniok%bas10.basnet.minsk.by@demos.su
Martina Kratzsch	inmkr@sun.rz.tu-clausthal.de
Marek Kubale	kubale@sunrise.pg.gda.pl
Wieslaw Kubiak	wkubiak@kean.ucs.mun.ca
Jan Karel Lenstra	jkl@win.tue.nl
Eric Luit	wsinluit@win.tue.nl
Rolf Möhring	moehring@math.tu-berlin.de
Jerzy Nawrocki	nawrocki@pozn1v.tup.edu.pl
Erwin Pesch	e.pesch@uni-bonn.de
Jürgen Sauer	sauer@informatik.uni-oldenburg.de
Günter Schmidt	gs@wi2.uni-sb.de
Philippe Solot	isolot@chbs.ciba.com
V. Sridharan	suhas@clemson.edu
Eugeniusz Toczylowski	e.toczylowski@ia.pw.edu.pl
Denis Trystram	trystram@imag.fr
Jan Weglarz	weglarz@pozn1v.tup.edu.pl
Frank Werner	frank.werner@mathematik.uni-magdeburg.d400.de
Jia Xu	jxu@cs.yorku.ca

Dagstuhl-Seminar 9520

List of Participants

Jörg **Apsel**
TU Clausthal
Institut für Informatik
Erzstr. 1
38678 Clausthal-Zellerfeld
Germany
apsel@informatik.tu-clausthal.de
tel.: +49-5323-72-35 74

Euripides **Bampis**
Université Paris-Sud
LRI
Bat. 490
F-91405 Orsay
France
bampis@lri.fr
tel.: +33-1-69.41.64.53

Jacek **Blazewicz**
Politechnika Poznanska
Instytut Informatyki
ul. Piotrowo 3A
PL-60-965 Poznan
Poland
blazewic@pozn1v.tup.edu.pl
tel.: +48-61 - 790 790

Jean-Paul **Boufflet**
Université de Technologie de Compiegne
Dept. Genie Informatique
B.P. 649
F-60206 Compiegne
France
boufflet@hds.univ-compieqne.fr
tel.: +33-44 23 46 91

Pascal **Bouvry**
CWI
Dept. of I. S.
Kruislaan 413
NL-1090 GB Amsterdam
The Netherlands
bouvry@cwi.nl
tel.: +31 20 592 4009

Peter **Brucker**
Universität Osnabrück
FB Mathematik/Informatik
49069 Osnabrück
Germany
pbrucker@dosuni1.rz.uni-osnabrueck.de

Rainer **Burkard**
TU Graz
Institut für Mathematik
Steyrergasse 30
A-8010 Graz
Austria
rburkard@ftug.dnet.tu-graz.ac.at
tel.: +43-316-8737125

Jean-Yves **Colin**
LITP
Université Paris VI
75252 Paris Cedex 05
France
colin@lachesis.ibp.fr

Yves **Crama**
Université de Liege
Departement de Gestion
Boulevard du Rectorat 7 (B31)
B-4000 Liege
Belgium
crama@vm1.ulg.ac.be
tel.: +32-41-66 30 77

Moshe **Dror**
University of Arizona
College of Business & Public Administration
MIS Department
Tucson AZ 85721
USA
mdror@bpa.arizona.edu
tel.: +1-602-621-79 42

Maciej **Drozdowski**
Politechnka Poznanska
Instytut Informatyki
ul. Piotrowo 3A
PL-60-965 Poznan
Poland
maciej
tel.: +48-61-782 372

Klaus H. **Ecker**
TU Clausthal
Institut für Informatik
Leibnitzstr. 19
38678 Clausthal-Zellerfeld
Germany
ecker@informatik.tu-clausthal.de
tel.: +49-5323-723553

Gerd **Finke**
Universite Joseph Fourier
Laboratoire ARTEMIS
B. P. 53
F-38041 Grenoble Cedex 9
France
gerd.finke@imag.fr
tel.: +33-76514298

Lucian **Finta**
INRIA
Centre Sophia Antipolis
2004 Route des Lucioles
F-06902 `Sophia Antipolis
France
Lucian.Finta@sophia.inria.fr
tel.: +33-93 65 76 53

Valery **Gordon**
Belarus Academy of Sciences
Insitute of Engineering Cybernetics
ul Surganova 6
220012 Minsk
Belarus
mahaniok%bas10.basnet.minsk.by@demos.su
tel.: +7-0172-39 51 25

Frederic **Guinand**
LMC-IMAG
46 avenue Felix Viallet
F-38031 Grenoble Cedex
France
guinand@imag.fr
tel.. +33-76574870

Jeet **Gupta**
Ball State University
Department of Management
Munncie IN 47306
USA
00jngupta@leo.bsuvc.bsu.edu
tel.: +1-317-285-5301

Wolfgang **Halang**
FernUniversität Hagen
FB Elektrotechnik
Lehrgebiet Informationstechnik
58084 Hagen
Germany
wolfgang.halang@fernuni-hagen.de
tel.: +49-2331-98 73 72

Claire **Hanen**
Université P. et M. Curie
LITP
Tour 55-65 E428
4 place Jussieu
75252 Paris Cedex 05
France
hanen@litp.ibp.fr
tel.: +33-1-44 27 70 91

Reiner **Hirschberg**
pdv-systeme
Bornhardt Str. 3
38644 Goslar
Germany
hirsch@zaphad.in.tu-clausthal.de
tel.: +49-5321-37030

Adam **Janiak**
Politechnika Wroclawska
Instytut Cybernetyki Technicznej
ul. Janiszewskiego 11/17
PL-50-372 Wroclaw
Poland
Janiak@ict.pwr.wroc.pl
tel.: +48-71-202745

Jörg **Kaiser**
GMD
Institut für Systementwurfstechnik (SET)
Bereich: Responsive Systeme (RS)
Postfach 13 16
53731 St. Augustin
Germany
kaiser@gmd.de
tel.: +49-2241-142323

Roger **Kieckhafer**
University of Nebraska
Dept. of Computer Science
Lincoln NE 68588-0115
USA
rogerk@cse.unl.edu
tel.: +1-402-472-24 01

Mikhail **Kovalyov**
Belarus Academy of Sciences
Insitute of Engineering Cybernetics
ul. Surganova 6
220012 Minsk
Belarus
msmikhai-at-msnw01po@smtpg-
wy.polyu.edu.hk
tel.: +7-0172-39 51 71

Martina **Kratzsch**
TU Clausthal
Institut für Informatik
Leibnitzstr. 19
38678 Clausthal-Zellerfeld
Germany
inmkr@sun.rz.tu-clausthal.de
tel.: +49-5323-8 37 01

Marek **Kubale**
Politechnika Gdanska
Katedra Podstaw Informatyki
ul. Narutowicza 11/12
PL-80-952 Gdansk
Poland
kubale@sunrise.pg.gda.pl
tel.: +48-58-471818

Wieslaw **Kubiak**
Memorial University of Newfoundland
Faculty of Business Administration
St. John`s Newfoundland A1B 3X5
Canada
wkubiak@kean.ucs.mun.ca
tel.: +1-709-737-4412

Jan Karel **Lenstra**
Eindhoven University of Technology
Dept.of Mathematics and
Computing Science
P.O. Box 513
NL-5600 MB Eindhoven
The Netherlands
jkl@win.tue.nl
tel.: +31-40-473130 (secret.)
+31-40-4747 70 (office)

Erik **Luit**
Eindhoven University of Technology
Dept.of Mathematics and
Computing Science
Den Dolech
NL-5600 MB Eindhoven
The Netherlands
wsinluit@win.tue.nl
tel.: +31-40-474-284

Rolf H. **Möhring**
TU Berlin
FB Mathematik
MA 6-1
Str. d. 17. Juni 136
10623 Berlin
Germany
moehring@math.tu-berlin.de
tel.: +49-30-314-24594

Jerzy R. **Nawrocki**
Politechnka Poznanska
Instytut Informatyki
ul. Piotrowo 3A
PL-60-965 Poznan
Poland
nawrocki@pozn1v.tup.edu.pl
tel.: +48 61 78 23 75

Erwin **Pesch**
Universität Bonn
Institut für Gesellschafts- und
Wirtschaftswissenschaften
Adenauerallee 24-42
D-53113 Bonn
Germany
e.pesch@uni-bonn.de
tel.: +49-228-739-2 47

Jürgen **Sauer**
Universität Oldenburg
FB 10 - Informatik
Ammerländer Heerstr. 114
D-26111 Oldenburg
Germany
sauer@informatik.uni-oldenburg.de
tel.: +49-441-9706-2 95

Günter **Schmidt**
Universität des Saarlandes
Postfach 15 11 50
66041 Saarbrücken
Germany
gs@wi2.uni-sb.de
tel.: +49-681-302-4559

Philippe **Solot**
Ciba-Geigy AG
Operations Research
IS 7.2 (R-1008.Z2.27)
Postfach
CH-4002 Basel
Switzerland
isolot@chbs.ciba.com
tel.: +41-61-6979-012

V. **Sridharan**
Clemson University
Department of Management
142 Sirrine Hall
Clemson SC 29634-1305
USA
suhas@clemson.edu
tel.: +1-803-656-2624

Eugeniusz **Toczylowski**
Politechnika Warszawska
Instytut Automatyki
Plac Jednosci Robotniczej
Warszawa
Poland
e.toczylowski@ia.pw.edu.pl
tel.: +48-2-660-79 50

Denis **Trystram**
LMC-IMAG
46 avenue Felix Viallet
F-38031 Grenoble Cedex
France
trystram@imag.fr
tel.: +33-76574876

Jan **Weglarz**
Politechnika Poznanska
Instytut Informatyki
ul. Piotrowo 3A
PL-60-965 Poznan
Poland
weglarz@pozn1v.tup.edu.pl
tel.: +48-61-790790

Frank **Werner**
Otto-von-Guericke-Universität
Fakultät für Mathematik
Institut für Mathematische Optimierung
Universitätsplatz 2
39106 Magdeburg
Germany
frank.werner@mathematik.uni-magde-
burg.d400.de
tel.: +49-391-5592-2025

Jia **Xu**
York University
Department of Computer Science
4700 Keele Street
North York Ontario M3J-1P3
Canada
jxu@cs.yorku.ca
tel.: +1-416-736-50 53

Zuletzt erschienene und geplante Titel:

W. Hoeppner, H. Horacek, J. Moore (editors):
Principles of Natural Language Generation, Dagstuhl-Seminar-Report; 93; 25.07.-29.07.94 (9430)

A. Lesgold, F. Schmalhofer (editors):
Expert- and Tutoring-Systems as Media for Embodying and Sharing Knowledge, Dagstuhl-Seminar-Report; 94; 01.08.-05.08.94 (9431)

H.-D. Ehrich, G. Engels, J. Paredaens, P. Wegner (editors):
Fundamentals of Object-Oriented Languages, Systems, and Methods, Dagstuhl-Seminar-Report; 95; 22.08.-26.08.94 (9434)

K. Birman, F. Cristian, F. Mattern, A. Schiper (editors):
Unifying Theory and Practice in Distributed Systems, Dagstuhl-Seminar-Report; 96; 05.09.-09.09.94 (9436)

R. Keil-Slawik, I. Wagner (editors):
Interdisciplinary Foundations of Systems Design and Evaluation, Dagstuhl-Seminar-Report; 97; 19.09.-23.09.94 (9438)

M. Broy, L. Lamport (editors):
Specification and Refinement of Reactive Systems - A Case Study, Dagstuhl-Seminar-Report; 98; 26.09.-30.09.94 (9439)

M. Jarke, P. Loucopoulos, J. Mylopoulos, A. Sutcliffe (editors):
System Requirements: Analysis, Management, and Exploitation, Dagstuhl-Seminar-Report; 99; 04.10.-07.10.94 (9440)

J. Buchmann, H. Niederreiter, A.M. Odlyzko, H.G. Zimmer (editors):
Algorithms and Number Theory, Dagstuhl-Seminar-Report; 100; 10.10.-14.10.94 (9441)

S. Heinrich, J. Traub, H. Wozniakowski (editors):
Algorithms and Complexity for Continuous Problems, Dagstuhl-Seminar-Report; 101; 17.10.-21.10.94 (9442)

H. Bunke, T. Kanade, H. Noltemeier (editors):
Environment Modelling and Motion Planning for Autonomous Robots, Dagstuhl-Seminar-Report; 102; 24.10.-28.10.94 (9443)

W. Maass, Ch. v.d. Malsburg, E. Sontag, I. Wegener (editors):
Neural Computing, Dagstuhl-Seminar-Report; 103; 07.11.-11.11.94 (9445)

G. Berry, W.P. de Roever, A. Poigné, A. Pnueli (editors):
Syncronous Languages, Dagstuhl-Seminar-Report; 104; 28.11.-02.12.94 (9448)

B. Becker, R. Bryant, O. Coudert, Ch. Meinel (editors):
Computer Aided Design and Test, Dagstuhl-Seminar-Report; 105; 13.02.-17.02.95 (9507)

D. Garlan, F. Paulisch, W. Tichy (editors):
Software Architectures, Dagstuhl-Seminar-Report; 106; 20.02.-24.02.95 (9508)

W.J. Cullyer, W.A. Halang, B. Krämer (editors):
High Integrity Programmable Electronic Systems, Dagstuhl-Seminar-Report; 107; 27.02.-03.03.95 (9509)

J. Gruska, T. Toffoli, H. Umeo, R. Vollmar (editors):
Cellular Automata, Dagstuhl-Seminar-Report; 108; 06.03.-10.03.95 (9510)

H. Alt, B. Chazelle, R. Seidel (editors):
Computational Geometry, Dagstuhl-Seminar-Report; 109; 13.03.-17.03.95 (9511)

W. Bibel, K. Furukawa, M. Stickel (editors):
Deduction, Dagstuhl-Seminar-Report; 110; 20.03.-24.03.95 (9512)

B. Freitag, C.B. Jones, Ch. Lengauer, H.-J. Scheck (editors):
Object-Orientation with Parallelism and Persistence, Dagstuhl-Seminar-Report; 111; 03.04.-07.04.95 (9514)

J. Doran, N. Gilbert, U. Mueller, K. Troitzsch (editors):
Social Science Microsimulation: A Challenge for Computer Science, Dagstuhl-Seminar-Report; 112; 01.05.-05.05.95 (9518)

R. Studer, M. Musen (editors):
Shareable and Reusable Problem Solving Methods, Dagstuhl-Seminar-Report; 113; 08.05.-12.05.95 (9519)

J. Blazewicz, K. Ecker, L. Welch (editors):
Scheduling in Computer & Manufacturing Systems, Dagstuhl-Seminar-Report; 114; 15.05.-19.05.95 (9520)

H. Beilner, G. Ciardo, C. Lindemann, K. Trivedi (editors):
Performance and Dependability Modeling with Stochastic Petri Nets, Dagstuhl-Seminar-Report; 115; 22.05.-26.05.95 (9521)

M. Aigner, J. Spencer, E. Triesch (editors):
Computing with Faulty Inputs, Dagstuhl-Seminar-Report; 116; 29.05.-02.06.95 (9522)

J.-R. Abrial, E. Börger, H. Langmaack (editors):
Methods for Semantics and Specification, Dagstuhl-Seminar-Report; 117; 05.06.-09.06.95 (9523)

W. Effelsberg, D. Ferrari, O. Spaniol, A. Danthine (editors):
Architecture and Protocols for High Performance Networks, Dagstuhl-Seminar-Report; 118; 19.06.-23.06.95 (9525)

Ph. Flajolet, R. Kemp, H. Prodinger, R. Sedgewick (editors):
`Average-Case'-Analysis of Algorithms, Dagstuhl-Seminar-Report; 119; 03.07.-07.07.95 (9527)

D. Gustfield, T. Lengauer, C. Sander (editors):
Molecular Bioinformatics, Dagstuhl-Seminar-Report; 120; 10.07.-14.07.95 (9528)

J. Chomicki, G. Saake, C. Sernadas (editors):
Role of Logics in Information Systems, Dagstuhl-Seminar-Report; 121; 17.07.-21.07.95 (9529)

R.S. Boyer,A. Bundy, D.Kapur, Ch. Walther (editors):
Automation of Proof by Mathematical Induction, Dagstuhl-Seminar-Report; 122; 24.07.-28.07.95 (9530)

P. Cousot, R. Cousot, A. Mycroft (editors):
Abstract Interpretation, Dagstuhl-Seminar-Report; 123; 28.08.-01.09.95 (9535)

P. Brunet, D. Roller, J. Rossignac (editors):
CAD Tools for Products, Dagstuhl-Seminar-Report; 124; 04.09.-08.09.95 (9536)

C. Dwork, E.W. Mayr, F. Meyer a.d. Heide (editors):
Parallel and Distributed Algorithms, Dagstuhl-Seminar-Report; 125; 11.09.-15.09.95 (9537)

C. Hankin, H. R. Nielson (editors):
New Trends In the Integration of Paradigms, Dagstuhl-Seminar-Report; 126; 18.09.-22.09.95 (9538)

U. Herzog, G. Latouche, P. Tran-Gia, V. Ramaswami (editors):
Applied Stochastic Modelling in Telecommunication and Manufacturing Systems, Dagstuhl-Seminar-Report; 127; 25.09.-29.09.95 (9539)

L. Hordijk, G. Korn, A. Sydow (editors):
Modelling and Simulation of Complex Environmental Problems, Dagstuhl-Seminar-Report; 128; 02.10.-06.10.95 (9540)