Dagstuhl Seminar on

Online Algorithms

July 1–5 2002

Organized by:

Susanne Albers (Freiburg University)
Amos Fiat (Tel-Aviv University)
Gerhard Woeginger (Twente University)

Summary

Online algorithms have received considerable research interest during the last 15 years. In an online problem the input arrives incrementally, one piece at a time. In response to each input portion, an online algorithm must generate output, not knowing future input. Online problems arise in very many areas of computer science, including e.g. resource allocation in operating systems, data structuring, robotics or large networks. The performance of online algorithms is usually evaluated using competitive analysis. An online algorithm A is called c-competitive if, for all input sequences, the solution computed by A is at most a factor of c away from the solution generated by an optimal offline algorithm that knows the entire input in advance.

The Dagstuhl meeting on *Online Algorithms* brought together 58 researchers with affiliations in Austria, the Czech Republic, Denmark, Germany, Hong Kong, Israel, Italy, Japan, the Netherlands, Switzerland, the UK and the USA. 11 participants were young scientists. There were 40 workshop presentations and, despite of the large number of talks, there was sufficient time for scientific discussions and research in teams. The presentations addressed classical online problems as well as new problems arising in important applications of current interest. Many interesting results were presented on classical issues such as paging and caching, bin packing, coloring, the k-server problem and metrical task systems. There was also a considerable number of lectures on online scheduling problems. Several presentations considered the fresh and interesting field of competitive auctions and game theory. With respect to application areas, many talks investigated problems that arise in large networks. Moreover there were talks studying problems in robotics, online learning, media-on-demand, power saving, seat reservation and vehicle routing. On Thursday evening there was an open problem session where interesting and new problems were presented.

Thanks: The abstracts of the 40 talks were compiled by Rob van Stee. The organizers would like to thank all participants and especially the team of Schloß Dagstuhl for helping to make the workshop a success. The warm atmosphere of the Schloß, as always, supported discussions and the informal exchange of ideas!

Abstracts

On Randomized Online Scheduling
Susanne Albers

We study one of the most basic problems in online scheduling. A sequence of jobs must be scheduled on m identical parallel machines so as to minimize the makespan. We present the first randomized online algorithm that beats known deterministic algorithms for general m. Our new algorithm is 1.916-competitive. We prove that this performance cannot be proven for a deterministic online algorithm based on the analysis techniques that have been used in the literature so far.

Auctions of Network Resources
Yossi Azar

We consider the problem of network routing and admission control under an auction model. Here requests for connections arrive online along with bids offered if the connection is routed. We design a truthtelling mechanism for this problem which obtains within a factor of $O(\log \mu)$ of the best possible profit for the network controller. Here μ is the product of the number of network nodes and the ratio of maximum to minimum bids. We present matching lower bounds to show that this result is best-possible (within constant factors). A modified version of this algorithm is resilient against coalitions and retries. The algorithm depends upon prior knowledge of the maximum possible bid. Without this knowledge, we present a competitive algorithm with a slightly larger competitive ratio along with near-matching lower bounds. These results will also hold in a general combinatorial auction where customers request a subset of the resources upon which to bid. All previous results for competitive algorithms for auction were provided only for a single (type of) resource.

Joint work with Baruch Awerbuch and Adam Meyerson.

Competitive On-Line Stream Merging Algorithms for Media-on-Demand
Amotz Bar-Noy

We consider the problem of minimizing the bandwidth needed by media-on-demand servers that use stream merging. We consider the on-line case where client requests are not known ahead of time. To facilitate stream merging, clients have the ability to receive data from two streams simultaneously and can buffer up to half of a full stream. We present a new family of on-line stream merging algorithms called *dynamic tree algorithms*. The bandwidth requirements of the best of these, the *dynamic Fibonacci tree algorithms*, are within a factor of the minimum between $\log_\phi(n)+O(1)$ and $\log_\phi(1/(2D))+O(1)$ from the off-line optimal, where n is the number of requests, D is the guaranteed maximum startup delay measured as a fraction of the time for a full stream, and $\phi = (1+\sqrt{5})/2$. The new on-line algorithms use a dynamic Fibonacci tree to control how new arrivals should merge with existing streams. Empirical studies show that the dynamic Fibonacci tree algorithms perform much better than indicated by the analysis.

Joint work with Richard E. Ladner.

[1] A. Bar-Noy and R. E. Ladner. Competitive On-Line Stream Merging Algorithms for Media-on-Demand. *the 12th ACM-SIAM Symp. on Discrete Algorithms (SODA)*, pp. 364–373, 2001.

Knowledge State Algorithms

WOLFGANG W. BEIN

We introduce the novel concept of knowledge states. A knowledge state simply states conditional obligations of an adversary, by fixing a work function, and gives a distribution for the algorithm. When a knowledge state algorithm receives a request, it then calculates one or more "subsequent" knowledge states, together with a probability of transition to each. The algorithm then uses randomization to select one of those subsequents to be the new knowledge state. Although the formal definition of knowledge state algorithms appears, as yet, in no publication, many well-known algorithms can be viewed as knowledge state algorithms.

We have used optimization techniques to construct a non-trivial knowledge state algorithm for the 2-server problem on the line with competitive ratio $\frac{71}{36} \approx 1.972$. As much as one avenue of investigation might be to further improve this result for the line, what we really envision is to use this technique to finally settle the question of whether there exists an online algorithm with competitive ratio better than 2 for general spaces.

We have also constructed an optimally competitive knowledge state algorithm for a cache management problem in a multiprocessor architecture (this work is jointly with R. Reischuk.) Here each processor P_i manages its own cache, but whenever a processor P_j, $j \neq i$, writes to a page which has a copy in the cache of P_i, that copy has to be invalidated (and effectively ejected from P_i.) Therefore, in a multiprocessor system, there can be two reasons for a page miss: A processor P_{i_0}, with cache size k, can experience a "miss" in the usual way, *i.e.*, because a data item needed by the processor was not among any of its k cached pages. Alternatively, the page may have been in the cache, but an invalidation caused by some other processor might have occurred. This we call an invalidation miss. This leads to an optimization model with a tantalizing information restriction: Given the sequence ρ of requested pages for processor P_{i_0} one considers ρ to be known in its entirety to P_{i_0}'s caching strategy; only invalidations are given online. Using the knowledge state approach, we have been able to prove a competitive ratio of $\frac{4}{3}$ for $k=2$. For $k=3$ we have obtained a lower bound of $\frac{3}{2}$ and we conjecture that it is possible to find a knowledge state algorithm that matches the lower bound. As much as the lower bound carries over for $k > 3$, the upper bound is open.

Joint work with Larry Larmore.

Online learning tools for Online Algorithms
Avrim Blum

In this survey talk we describe a few results from online learning and show how they can be useful tools for designing online algorithms. The basic tools are (1) the "randomized weighted majority" algorithm for combining expert advice in the game-theoretic setting, (2) extensions that address movement costs in the oblivious adversary model, and (3) extensions to "bandit" versions of these problems. In essence, these tools allow one to combine multiple online algorithms in many natural settings. We show how these tools can be immediately applied to a number of basic online algorithm problems such as "picking a winner", adaptive data structures, metrical task systems, and the design of online auctions.

Portions of this are joint work with Carl Burch, Shuchi Chawla, and Adam Kalai.

The Seat Reservation Problem Allowing Seat Changes
Joan Boyar

We consider a variant of the Seat Reservation Problem in which seat changes are allowed. A very promising natural class of algorithms, *conservative algorithms*, is introduced. With these algorithms, a passenger is only asked to change seats if there is another passenger assigned to that seat, starting at the station where the seat change occurs. These algorithms all have very high performance guarantees. For instance, if the optimal off-line algorithm can seat all of the passengers, $\frac{2}{3}$ of the passengers can be seated on-line using any conservative algorithm allowing only one seat change and $\frac{3}{4}$ will be seated if two seat changes are allowed. This should be compared to the asymptotic hardness result of $\frac{1}{2}$ for the best algorithm when no seat changes are allowed. In addition, a new algorithm, Modified-Kierstead-Trotter, is proposed for the case in which one seat change is allowed and shown to seat all passengers if the optimal off-line algorithm could have accommodated them with only half as many seats.

Joint work with Susan Krarup and Morten Nyhave Nielsen.

Algorithms for On-Line Bin-Packing Problems with Cardinality Constraints
Bo Chen

The bin packing problem asks for a packing of a list of items of sizes from $(0, 1]$ into the smallest possible number of bins having unit capacity. The k-item bin packing problem additionally imposes the constraint that at most k items are allowed in one bin. We present two efficient on-line algorithms for this problem. We show that, for increasing values of k, the bound on the asymptotic worst-case performance ratio of the first algorithm tends towards 2 and that the second algorithm has a Ratio of 2. Both algorithms considerably improve upon the best known result of an algorithm, which has an asymptotic bound of 2.7 on its Ratio. Moreover, we improve known bounds for all values of k by presenting on-line algorithms for $k = 2$ and $k = 3$ with bounds on their Ratios close to optimal.

Joint work with Luitpold Babel, Hans Kellerer and Vladimir Kotov.

Towards the Analysis of the Harmonic Algorithm for k-Servers

MAREK CHROBAK

Techniques from electrical network theory have been used to establish various properties of random walks. We explore this connection further, by showing how the classical formulas for the determinant and cofactors of the admittance matrix, due to Maxwell and Kirchoff, yield upper bounds on the edge stretch factor of the harmonic random walk. For any complete, n-vertex graph with distances assigned to its edges, we show the upper bound of $(n-1)^2$. If the distance function satisfies the triangle inequality, we give the upper bound of $\frac{1}{2}n(n-1)$. Both bounds are tight.

As a consequence, we obtain some new results on the competitive ratio of the the harmonic algorithm for the k server problem. HARMONIC is a randomized k-server algorithm that, at each step, given a request point r, chooses the server to be moved to r with probability inversely proportional to the distance to r. For general k, it is known that the competitive ratio of HARMONIC is at least $\frac{1}{2}k(k+1)$, while the best upper bound on this ratio is exponential in k. It has been conjectured that the lower bound is tight.

We provide further evidence for this conjecture, by proving two results. The first result states that HARMONIC is $\frac{1}{2}k(k+1)$-competitive against the lazy adversary, that is, the adversary that moves a server only when all his servers covered by our servers.

Our second result is that HARMONIC is 6-competitive for $k = 3$. We propose a new potential function Φ (for arbitrary k) and reduce the proof of the validity of Φ to several inequalities involving hitting costs of harmonic random walks. Then, we show that these inequalities hold for $k = 3$.

Joint with with Yair Bartal, John Noga, Prabhakar Raghavan and Jiří Sgall.

Minimizing the maximum starting time on-line

LEAH EPSTEIN

We study the scheduling problem of minimizing the maximum starting time on-line. The goal is to minimize the last time that a job starts. We show that while the greedy algorithm has a competitive ratio of $O(\log m)$, we can give a constant competitive algorithm for this problem. We also show that the greedy algorithm is optimal for resource augmentation in the sense that it requires $2m-1$ machines to have a competitive ratio of 1, whereas no algorithm can achieve this with $2m-2$ machines.

Joint work with Rob van Stee.

On Paging with Locality of Reference

Lene Monrad Favhroldt

Motivated by the fact that competitive analysis yields too pessimistic results when applied to the paging problem, there has been considerable research interest in refining competitive analysis and in developing alternative models for studying online paging. The goal is to devise models in which theoretical results capture phenomena observed in practice.

In this paper we propose a new, simple model for studying paging with locality of reference. The model is closely related to Denning's working set concept and directly reflects the amount of locality that request sequences exhibit. We demonstrate that our model is reasonable from a practical point of view.

We use the page fault rate to evaluate the quality of paging algorithms, which is the performance measure used in practice. We develop tight or nearly tight bounds on the fault rates achieved by popular paging algorithms such as LRU, FIFO, deterministic Marking strategies and LFD. It shows that LRU is an optimal online algorithm, whereas FIFO and Marking strategies are not optimal in general. We present an experimental study comparing the page fault rates proven in our analyses to the page fault rates observed in practice. This is the first such study for an alternative/refined paging model.

Joint work with Susanne Albers and Oliver Giel.

Companion Caching

Amos Fiat

Steve Seiden died in a tragic accident on June 11, 2002.

This paper is concerned with online caching algorithms for the (n,k)-companion cache, defined by Brehob et al. [1]. In this model the cache is composed of two components: a k-way set-associative cache and a companion fully-associative cache of size n. We show that the deterministic competitive ratio for this problem is $(n+1)(k+1)-1$, and the randomized competitive ratio is $O(\log n \log k)$ and $\Omega(\log n + \log k)$.

Joint work with Manor Mendel and Steve Seiden.

[1] M. Brehob, R. Enbody, E. Torng, and S. Wagner. On-line restricted caching. In *Proceedings of the 12th Symposium on Discrete Algorithms*, pages 374–383, 2001.

MRE Routing

Rudolf Fleischer

We study Maximum Residual Energy (MRE) routing in battery powered wireless networks where nodes need energy to send and receive messages (because messages must be acknowledged). This models certain Bluetooth applications, for example. We show that the general problem is NP-complete and cannot be approximated by an exponential factor.

Joint work with Mordecai Golin, Chin-Tau Lea, and Steven Wong.

New results on competitive auctions
ANDREW GOLDBERG

We introduce the following consensus estimate problem. Several processors hold private lower bounds on a value. The processors do not communicate with each other, but can observe a shared source of random numbers. The goal is to come up with a consensus lower bound on the value that is as high as possible. We give a solution to the consensus estimate problem and show how to apply this solution in the context of mechanism design.

Based on our consensus estimate technique, we introduce Consensus Revenue Estimate (CORE) auctions. This is a class of competitive auctions that is interesting for several reasons. One auction from this class achieves a better competitive ratio than any previously known auction. Another one uses only two random bits, whereas the previously known competitive auctions on n bidders use n random bits. A parameterized CORE auction performs better than the previous auctions in the context of mass-market goods, such as digital goods.

We also study the double auction (market clearing) problem. We develop a framework for designing competitive auctions for this problem, which is more general than the basic auction problem, and give several competitive auctions.

Joint work with Kaustubh Deshmukh, Jason Hartline and Anna Karlin.

Two level online problems
CSANÁD IMREH

Many interesting problem arise from well-investigated online problems if we give the possibility of making some extra decision to the algorithms. In the talk we gave a short overview about this kind of problems. We presented three classes of such problems.

- In the first class we suppose that the amount of resources (which in the original model is a parameter of the problem) is not fixed the algorithm has to purchase it. Here we detailed the scheduling problem where it is allowed to buy new machines and the paging problem where it is allowed to buy extra cache.

- In the second class it is allowed to change the input. Here we detailed a strip packing problem where it is possible to strech the rectangles keeping the area fixed.

- In the third class the algorithm can choose between two problems with different parameters. Here we showed a two sets machine scheduling model, where we have two sets both containing identical machines, and the jobs have two different processing times one for each set of mechines.

Online Strategies for Power Savings
Sandy Irani

This talk is a survey of recent work in power saving strategies. We address two different mechanisms for saving power. The first is the ability to transition a device or system into a sleep state if it is idle. The trade-off is that a fixed amount of energy is needed to transition the device back into an active state in which it can begin to process jobs. We extend previous work on this problem and develop online algorithms for devices with multiple levels of sleep states. Each sleep state has a different power consumption level and different cost in transitioning back to the active state. The second way in which power savings can be achieved is by varying the speed at which jobs are run. We utilize a power consumption curve which indicates the power consumption level given a particular speed. The problem is to schedule arriving jobs in a way that minimizes total energy use so that each job is completed by its deadline. Our work addresses systems whose speed can be varied and which have a sleep state. We devise both offline and online algorithms for this problem.

Removable On-Line Knapsack Problems
Kazuo Iwama

We introduce an on-line model for a class of hand-making games such as Rummy and Mah-Jang. An input is a sequence of items, $u_1, \cdots, u_i, \cdots$ such that $0 < |u_i| \leq 1.0$. When u_i is given, the on-line player puts it into the bin and can discard any selected items currently in the bin (including u_i) under the condition that the total size of the remaining items is at most one. The goal is to make this total size as close to 1.0 as possible when the game ends. We also discuss the multi-bin model, where the player can select a bin out of the k ones which u_i is put into. We prove tight bounds for the competitive ratio of this problem, both for $k = 1$ and $k \geq 2$.

Joint work with Shiro Taketomi.

Dealing with Mobility, Dynamic Service Selection and Pricing
Bala Kalyanasundaram

Bandwidth limitations, resource greedy applications and an increasing number of users are straining wireless networks. Hence, algorithms to manage wireless bandwidth are critical. Optimality of overall system behavior can be achieved by allowing users to change service level and/or service provider for a (small) price. The ability to dynamically re-negotiate service gives the user the power to control QoS as well as minimize cost. On the other hand, the ability to change the price of a service level dynamically allows the service providers a better way to manage traffic, improve resource usage and most importantly maximize their profit. This can then provide a surprising win-win situation for BOTH the service providers AND the users.

In this paper, we present an easy to implement online algorithm to minimize the overall cost of individual mobile users. This online algorithm continuously receives pricing information and minimum QoS requirements. The algorithm then determines appropriate service level, service provider and time for re-negotiation dynamically so that it minimizes the overall cost for the user and allow the service provider to control capacity to maximize revenue. Our algorithm can handle many practical issues such as capacity limitations, arbitrary price fluctuations and loss/gain of service providers due to mobility. Our results do not assume any specific wireless technologies and can be applied to any environment that can employ dynamic pricing, including wired networks.

Joint work with John Waclawsky and Mahe Velauthapillai

Competitive Generalized Auctions
ANNA R. KARLIN

The emergence of the Internet as one of the most important arenas for resource sharing between parties with diverse and selfish interests has led to a number of fascinating and new algorithmic problems. In these problems, one must solicit the inputs to each computation from the participants whose goal is to manipulate the computation to their own advantage. We consider a class of dynamic pricing problems, the most basic of which is an auction, and discuss the design of truthful mechanisms, in which the selfish participants are motivated to reveal their true input. We present a number of results on the design of profit-maximizing truthful generalized auctions.

This is joint work with Amos Fiat, Andrew Goldberg and Jason Hartline.

Convergence Time to Nash Equilibria
ALEXANDER KESSELMAN

In this paper we study the number of steps required to reach a Nash Equilibrium. We consider a variety of problems related to on-line load balancing where each job behaves selfishly and the load is determined by either a weighted load function or the objective function of a congestion game. Such scenarios arise, in particular, when users communicate over the Internet acting according to their own interests, without regard to overall network performance. The main research efforts in Game Theory have mostly been concentrated on the existence and uniqueness of Nash Equilibrium while the main research efforts in Computer Science have been focused on finding a global optimum and the ratio between the Nash Equilibrium and the optimum. In this paper we design efficient strategies that converge to a Nash Equilibrium in a polynomial number of steps.

Joint work with Yishai Mansour.

Game Chromatic Number

HAL KIERSTEAD

The *coloring game* is played by two players, Alice and Bob, on a graph $G(V,E)$ with a set of colors X. Alice starts the game and then the players take turns playing. On each turn a player selects an uncolored vertex and legally colors it with a color from X. In the usual game a color α is legal for a vertex v if none of the neighbors of v has been previously colored with α. Other definitions of legal lead to other versions of the game. Alice wins the game if eventually all the vertices are legally colored. Bob wins the game if there comes a time when some uncolored vertex cannot be legally colored. The *game chromatic number*, $\chi_g(G)$, of G is the least t such that Alice has a winning strategy when $|X| = t$. We shall survey recent results on game chromatic number and its variants. Here we give some sample results for planar graphs.

Theorem 1 (Kierstead) *All planar graphs satisfy* $\chi_g(G) \leq 18$.

The *oriented coloring game* is played in a similar fashion, but X is replaced by a tournament T. A color α is legal for a vertex v if all colored neighbors w of v satisfy $v \to w$ iff $\alpha \to \beta$, where β is the color of w. (For technical reasons, if there is a directed path of length two between two vertices, their colors must also be different.) The *oriented game chromatic number*, $\operatorname{ogcn}(G)$, of G is the least t such that Alice has a winning strategy in the oriented coloring game when $|T| = t$.

Theorem 2 (Kierstead and Trotter) *There exists an integer c such that all planar graphs satisfy* $\operatorname{ogcn}(G) < c$.

The *d-relaxed coloring game* is played in the normal way except that a color α is legal for a vertex v if v has at most d neighbors already colored with α. The *d-relaxed game chromatic number*, $\chi_g^d(G)$, of G is the least t such that Alice has a winning strategy in the d-relaxed coloring game when $|X| = t$.

Theorem 3 (Dunn and Kierstead) *There exists an integer d $(= 132)$ such that all planar graphs satisfy* $\chi_g^d(G) \leq 3$.

The (a,b)-coloring game is played in the normal way except on each turn Alice colors a vertices and Bob colors b vertices. The *(a,b)-game chromatic number*, (a,b)-$\chi_g(G)$, of G is the least t such that Alice has a winning strategy in the (a,b)-coloring game when $|X| = t$.

Theorem 4 (Kierstead and Yang) *All planar graphs G satisfy* $(3,1)\text{-}\chi_g(G) \leq 8$.

Benefit task systems with applications to online allocation of servers in a web server farm
Tracy Kimbrel

Motivated by a practical problem of allocating servers over time in a so-called server farm environment, we define the following general Benefit Task System problem. Given for each time t a set of states S_t and a benefit function B_t from $S_t \times S_{t+1}$ to the nonnegative reals, find a sequence of states $\{s_t\}$ such that for each t, $s_t \in S_t$, with the goal of maximizing $\sum_t B_t(s_t, s_{t+1})$. We say that an algorithm uses lookahead L if s_t is a function of B_τ for $\tau \leq t+L$. It is easily shown that there are no competitive algorithms for $L = 0$. We give tight upper and lower bounds of $1 + 1/L$ on the competitive ratio of randomized algorithms, where we define the competitive ratio as the ratio of the optimal offline benefit to the online benefit. Our randomized algorithm requires only $O(\log L)$ random bits. We demonstrate that the deterministic problem is harder than the randomized problem by giving a deterministic lower bound of $1 + 1/L + 1/(L^2 + 1)$. We give two deterministic algorithms. The first has a better asymptotic competitive ratio of $1 + 4/(L - 7)$, and is based on the strategy of the randomized algorithm. This strategy is to forgo all benefit for one step, in order to move to a good state from which to follow an optimal path for some number of steps before repeating the process. The second deterministic algorithm, with competitive ratio $(1+1/L)(L+1)^{1/L}$ which is better for small L, is based on geometrically discounting the projected benefit on future time steps up to the lookahead distance L, and choosing an optimal path of length L based on these weighted benefits. The first step only of this path is taken, and then the process is repeated. Finally, for the special case $L = 1$, we give a lower bound of 4 which matches the upper bound of the future discounting algorithm.

Joint work with T.S. Jayram, Robert Krauthgamer, Baruch Schieber, and Maxim Sviridenko

Non-Abusiveness Helps: An $\mathcal{O}(1)$-Competitive Algorithm for Minimizing the Maximum Flow Time in the Online Traveling Salesman Problem
Sven O. Krumke

In the online traveling salesman problem (OLTSP) requests for visits to cities arrive online while the salesman is traveling. We study the $F_{\max}$ − OLTSP where the objective is to minimize the maximum flow time. This objective is particularly interesting for applications. For instance, it can be identified with the maximal dissatisfaction of customers. Alas, there can be no competitive algorithm, neither deterministic nor randomized. Moreover, in contrast to scheduling resource augmentation, e.g. providing the online algorithm with a faster server, does not help, the crucial difference being that servers move in space.

Hence, competitive analysis fails to distinguish online algorithms. This unsatisfactory situation motivates the search for alternative analysis methods.

We introduce a natural restriction on the adversary for the $F_{\max} - \text{OLTSP}$ on the real line. A *non-abusive adversary* may only move in a direction if there are yet unserved requests on this side. Our main result is an algorithm which achieves a constant competitive ratio against the non-abusive adversary.

Joint work with Luigi Laura, Maarten Lipmann, Alberto Marchetti-Spaccamela, Willem E. de Paepe, Diana Poensgen and Leen Stougie.

Optimally Competitive Algorithms for Uniform Metrical Task Systems
Larry Larmore

We give new upper bounds on the randomized competitiveness of the k-state uniform metrical task system, and discuss how linear programming can be used to obtain lower bounds for metrical task systems for arbitrary finite metric spaces.

For any $k \geq 2$, there is a $(\frac{3}{2}H_k - \frac{1}{2k})$ - competitive randomized stable distribution algorithm for the uniform metrical task system with k states. For $k = 2$, this result was known by Chrobak and Noga, and is best possible.

The upper bound technique can be refined to yield better upper bounds as well as lower bounds. For example, for $k = 3$, preliminary calculations indicate that the competitive ratio C for uniform metrical task systems is bounded by $2.527 \leq C \leq 2.545$. We conjecture that, for fixed k, these refinements will yield bounds which converge to the true competitiveness.

Joint work with Wolfgang Bein.

Packet Bundling
Kim S. Larsen

When messages, which are to be sent point-to-point in a network, become available at irregular intervals, a decision must be made each time a new message becomes available as to whether it should be sent immediately or if it is better to wait for more messages and send them all together. Because of physical properties of the networks, a certain minimum amount of time must elapse in between the transmission of two packets. Thus, whereas waiting delays the transmission of the current data, sending immediately may delay the transmission of the next data to become available even more.

We consider deterministic and randomized algorithms for this on-line problem, and characterize these by tight results under a new quality measure. It is interesting to note that our results are quite different from earlier work on the problem where the physical properties of the networks were emphasized less.

Joint work with Jens S. Frederiksen.

Parallel Scheduling Problems in CDMA Wireless Networks
Stefano Leonardi

Next generation 3G/4G wireless data networks allow multiple codes (or channels) to be allocated to a single user, where each code can support multiple data rates. Providing fine-grained QoS to users in such networks poses the two dimensional challenge of assigning both power (rate) and codes for every user. This gives rise to a new class of parallel scheduling problems. We abstract general downlink scheduling problems suitable for proposed next generation wireless data systems. This includes a communication-theoretic model for multirate wireless channels. In addition, while conventional focus has been on throughput maximization, we attempt to optimize the maximum response time of jobs, which is more suitable for stream of user requests. We present provable results on the algorithmic complexity of these scheduling problems. In particular, we are able to provide very simple, online algorithms for approximating the optimal maximum response time. This relies on resource augmented competitive analysis. We also perform an experimental study with realistic data of channel conditions and user requests to show that our algorithms are more accurate than our worst case analysis shows, and they provide fine-grained QoS to users effectively.

Joint work with Luca Becchetti, Suhas Digavi, Alberto Marchetti-Spaccamela, S. Muthukrishnan, Thiaga Nandagopal and Andrea Vitaletti.

Randomized Online Algorithms and Restricted Adversaries
Vincenzo Liberatore

Yao's principle is a fundamental techique for proving lower bounds on randomized algorithm and is based on a game-theoretical duality result by von Neumann. In this paper, we prove an extension of the principle to the case when one of the players is further constrained by a set of linear inequalities. The corresponding duality result is interpreted in a variety of algorithmic contexts, including multi-objective optimization problems, performance tail of randomized algorithms, constrained adversaries, resource augmentation method, smoothed analysis, high-probability results, and loose-competitiveness.

Buffer policies for QoS switches
Yishay Mansour

In this talk we overview a recent research direction of developing buffer policies that maximize the benefit. Our model has a single FIFO buffer of bounded capacity, and arriving packets may have different values. The aim of the buffer policy is to maximize the value of the packets sent from the buffer.

Control Message Aggregation in Group Communication Protocols
Seffi Naor

Reliable data transmission protocols between a sender and a receiver often use feedback from receiver to sender to acknowledge correct data delivery. Such feedback is typically sent as control messages by receiver nodes. Since sending of control messages involves communication overhead, many protocols rely on aggregating a number of control messages and sending them together as a single packet over the network. On the other hand, the delays in the transmission of control messages may reduce the rate of data transmission from the sender. Thus, there is a basic tradeoff between the communication cost of control messages and the effect of delaying them.

We develop a rigorous framework to study the aggregation of control packets for multicast and other hierarchical network protocols. We define the multicast aggregation problem and design efficient online algorithms for it, both centralized and distributed.

Joint work with Sanjeev Khanna and Dan Raz.

On-line dial-a-ride problems under a restricted information model
Willem de Paepe

In on-line dial-a-ride problems, servers are traveling in some metric space to serve requests for rides which are presented over time. Each ride is characterized by two points in the metric space, a *source*, the starting point of the ride, and a *destination*, the end point of the ride. Usually it is assumed that at the release of such a request complete information about the ride is known. We diverge from this by assuming that at the release of such a ride only information about the source is given. At visiting the source, the information about the destination will be made available to the servers. For many practical problems, our model is closer to reality. However, we feel that the lack of information is often a *choice*, rather than inherent to the problem: additional information *can* be obtained, but this requires investments in information systems. We give mathematical evidence that for the problem under study it pays to invest.

Joint work with Maarten Lipmann, X. Lu, René A. Sitters, and Leen Stougie.

Online Call Admission in Optical Networks with Larger Demands
Diana Poensgen

In the problem of *Online Call Admission in Optical Networks*, briefly called Oca, we are given a graph $G = (V, E)$ together with a set of wavelengths W and a finite sequence $\sigma = r_1, r_2, \ldots, r_m$ of calls which arrive in an online fashion. Each call r_j specifies a pair of nodes to be connected and an integral demand indicating the number of required lightpaths. A lightpath is a path in G together with a wavelength $\lambda \in W$.

Upon arrival of a call, an online algorithm must decide immediately and irrevocably whether to accept or to reject the call without any knowledge of calls which appear later in the sequence. If the call is accepted, the algorithm must provide the requested number of lightpaths to connect the specified nodes. The essential restriction is the wavelength conflict constraint: each wavelength is available only once per edge, which implies that two lightpaths sharing an edge must have different wavelengths. Each accepted call contributes a benefit equal to its demand to the overall profit. The objective in OCA is to maximize the overall profit.

Competitive algorithms for OCA have been known for the special case where every call requests just a single lightpath. In this paper we present the first competitive online algorithms for the general case of larger demands.

Joint work with Sven Krumke.

On-Line Scheduling of a Single Machine to Minimize Total Weighted Completion Time
CHRIS N. POTTS

We consider the on-line scheduling of a single machine in which jobs arrive over time, and preemption is not allowed. The goal is to minimize the total weighted completion time. We show that a simple modification of the shortest weighted processing time rule has a competitive ratio of 2. This result is established using a new proof technique which does not rely explicitly on a lower bound on the optimal objective function value. Since it is known that no on-line algorithm can have a competitive ratio of less than 2, we have resolved the open issue of determining the minimum competitive ratio for this problem

Joint work with Edward J. Anderson.

$(1+\epsilon)$-speed $O(1)$-competitive server scheduling algorithms
KIRK PRUHS

We argue that an analysis showing that a server scheduling algorithm is $(1+\epsilon)$-speed $O(1)$-competitive should be seen as strong evidence that this algorithm will perform well in practice. We use as examples nonclairvoyant CPU scheduling to minimize average flow time and multicast pull scheduling to minimize average flow time.

This talk contains joint work with Bala Kalyanasundaram and Jeff Edmonds.

Dynamic Routing on Networks with Fixed-Size Buffers
ADI ROSEN

The combination of the buffer sizes of routers deployed in the Internet and the Internet traffic itself leads routinely to routers dropping packets. Motivated by this, we initiate the rigorous study, using competitive analysis, of dynamic store-and-forward protocols on arbitrary networks in a setting where dropped packets are explicitly taken into account.

We assume that arbitrary traffic can be injected into the network, and that routers have buffers of fixed size, independent of network parameters, at the tail of each edge. We then analyze and compare the effectiveness of several online, local-control, greedy protocols and several network topologies using the competitive ratio of the throughput. Among other, we show that the protocol Nearest-To-Go (NTG) is competitive on all topologies, while the protocol Furthest-To-Go (FTG) is not competitive on certain topologies. Furthermore, on the specific topology of the line, NTG has a better competitive ratio than FTG.

These results are in sharp contrast to results obtained in the framework of Adversarial Queuing Theory (Borodin et al. 1996), where FTG is known to be stable on any topology while NTG can be unstable on certain topologies even at arbitrary low traffic rates. This suggests that stability analysis may not be the right means for comparing protocols on scalable networks with fixed-size buffers and arbitrary traffic.

Joint work with William Aiello, Eyal Kushilevitz, and Rafail Ostrovsky.

Continual Optimization, What, Why and How
BARUCH SCHIEBER

We use the term "continual optimization" to refer to the application of advanced analytic methods, such as mathematical optimization and special purpose heuristics, on combinations of real-time and historical data, to update business plans and make decisions that are communicated in real time for the purpose of optimizing business objectives. We now have an unprecedented opportunity to analyze and model complex systems with increasing accuracy and precision due to dramatic growth in the availability of processing power, the availability of robust software libraries, increases in the availability of real time digital data and fast, high-bandwidth communications capability. Most importantly, companies now need to use analytic tools as they engage in ever more competitive e-business.

In this talk we will discuss the evolution of continual optimization, describe related projects ongoing at IBM Research, and pose several research challenges motivated by it.

Online Scheduling of Jobs with Unknown Processing Times
UWE SCHWIEGELSHOHN

In this talk we assumed m identical machines on which independent and possibly parallel jobs must be executed. Each job j is characterized by its parallelity $m_j \leq m$ and a weight w_j. The jobs are submitted at different release dates r_j and are unknown before their release. Further, the processing time p_j is unknown until the job has completed. Jobs are assigned to machines once enough resources are available. Previous makespan results were briefly mentioned. It was further shown that for total weighted completion time no algorithm with bounded competitiveness exists even if all weights are 1. This result holds for sequential jobs with and without preemption. The same is true in the non-preemptive case if the ratio $\frac{w_j}{p_j}$ is given for each job at its release date. In this talk we addressed the so called sequential *equal ratio* case where all jobs require a single machine ($m_j = 1$) and $w_j = p_j$. We showed that the list scheduling algorithm is at least 1.250 competitive and conjectured that it achieves the lower bound of 1.207 given by Kawaguchi and Kyan.

To this end the main proof concepts were described. Finally, we presented some open questions with a relationship to this problem.

Joint work with Mohamed Hussein.

Preemptive Scheduling in Overloaded Systems
JIŘÍ SGALL

The following scheduling problem is studied: We are given a set of tasks with release times, deadlines, and profit rates. The objective is to determine a 1-processor preemptive schedule of the given tasks that maximizes the overall profit. In the standard model, each completed task brings profit, while non-completed tasks do not. In the metered model, a task brings profit proportional to the execution time even if not completed. For the metered task model, we present an efficient offline algorithm and improve both the lower and upper bounds on the competitive ratio of online algorithms. Furthermore, we prove three lower bound results concerning resource augmentation in both models.

Joint work with Marek Chrobak, Leah Epstein, John Noga, Rob van Stee, Tomáš Tichý and Nodari Vakhania.

New Bounds for Variable-Sized and Resource Augmented Online Bin Packing
ROB VAN STEE

In the variable-sized online bin packing problem, one has to assign items to bins one by one. The bins are drawn from some fixed set of sizes, and the goal is to minimize the sum of the sizes of the bins used. We present new algorithms for this problem and show upper bounds for them which improve on the best previous upper bounds. We also show the first general lower bounds for this problem. The case where bins of two sizes, 1 and $\alpha \in (0, 1)$, are used is studied in detail. This investigation leads us to the discovery of several interesting fractal-like curves. Our techniques are also applicable to the closely related resource augmented online bin packing problem, where we have also obtained the first general lower bounds.

Joint work with Steve Seiden and Leah Epstein.

We dedicate this paper to the memory of Steve.

Randomized On-line Scheduling on 3 processors
TOMÁŠ TICHÝ

We study randomized on-line non-preemptive scheduling in multiprocessor systems. In this problem each task is specified by its processing time and scheduled on any of m identical processors. The objective is to minimize the competitive ratio, which is the maximum of the ratio of the expected makespan and the optimal makespan over all input sequences.

The main result is an improvement on the lower bound on the competitive ratio for $m = 3$ processors: we show the competitive ratio must be strictly greater than 27/19.

An Exact Resource Augmentation Analysis of the SRPT Algorithm
Eric Torng

In this talk, we consider the problem of minimizing total flow time on a set of m identical machines where preemption is allowed and jobs are sequential and have release dates. We focus on the well-known Shortest Remaining Processing Time (SRPT) algorithm. We use the resource augmentation technique popularized by Kalyanasundaram and Pruhs to analyze SRPT. We improve upon a previous result by Phillips, Stein, Torng, and Wein to show that SRPT, when given processors that are at least $s \geq 2 - 1/m$ times as fast as those given to the optimal offline algorithm, incurs a flow time that is no more than $1/s$ times the flow time of the optimal offline algorithm. That is, SRPT is an s-speed $1/s$-approximation algorithm for this problem. This result is tight as can be seen by any input instance with a single job.

Joint work with Jason McCullough.

Experimental Studies of Graph Traversal Algorithms
Gerhard Wolfgang Trippen

In this paper we conduct an extensive experimental evaluation of graph traversal algorithms. We also suggest changes in existing algorithms motivated by observations made during the visualization of the algorithms and interactive experimentations with them. Several natural and more sophisticated ways of graph traversals are studied. Our work helps to provide a better insight into the practical performance by testing different graph families.

Joint work with Rudolf Fleischer.

Open Problems