

Dagstuhl Seminar No. 02251, Report No. 345
June 16 - 21 2002

Approximation and Randomized Algorithms in Communication Networks

E. Bampis (Evry, France)

K. Jansen (Kiel, Germany)

G. Persiano (Salerno, Italy)

R. Solis-Oba (London, Canada)

G. Wilfong (Bell Labs, Murray Hill, USA)

1 Overview

During the week of June 16 - 21, 2002, the seminar on Approximation and Randomized Algorithms in Communication Networks was organized by E. Bampis (Evry, France), K. Jansen (Kiel, Germany), G. Persiano (Salerno, Italy), R. Solis-Oba (London, Canada), G. Wilfong (Bell Labs, Murray Hill, USA). 45 Participants came from universities or research institutes from Canada, Cyprus, Greece, France, Germany, Israel, Italy, Netherlands, Switzerland, United Kingdom and United States of America.

The recent progress in network technologies and availability of large distributed computer systems has increased the need for efficient algorithms for solving the diverse optimization problems that arise in the management and usage of communication networks. Technological developments in communication networks, like broad-band, all-optical, and ATM networks have made this area very interesting and important in recent years. They have also created new research directions and projects. The objectives of this seminar are of both theoretical and practical significance. The seminar aims to contribute to the theory of approximation, randomized, and on-line algorithms for problems arising in communication networks. It also has as a goal to explore the use of this theory in the solution of real world applications and in the development of practical algorithmic tools, thus fostering the cooperation among theoretical and practical researchers in this field.

The topics of the seminar included: routing and communication in networks, design of high performance networks, wavelength routing in optical networks, ATM network problems, quality of service, robustness issues, frequency assignment in radio networks, time and resource constrained scheduling, scheduling with communication delays, load balancing, and resource allocation.

The seminar was intended to bring together researchers from different areas in combinatorial optimization and from applications. It would support the collaboration between researchers in Computer Science, Engineering, Mathematics, and related areas. Different algorithmic methods and techniques have been covered by 31 lectures.

The seminar had the following goals:

- pose new optimization problems arising from applications in communication networks,
- design improved approximation algorithms for optimization problems in communication networks,
- study new algorithmic methods using randomization, linear, and nonlinear programming,

- discuss the practical implementation of different techniques and methods proposed for solving network communication problems,
- exchange information on recent research and stimulate further research in this area.

Schloß Dagstuhl and its staff provided a very convenient and stimulating atmosphere. The organizers wish to thank the local organization and all those who helped to make the seminar a fruitful research experience.

On behalf of the participants,

Evripidis Bampis, Klaus Jansen, Giuseppe Persiano, Roberto Solis-Oba, and Gordon Wilfong

2 Program

Monday Morning Session Chair: Klaus Jansen

9:15 *Welcome*

9:30 - 10:00 *Tight Approximation results for general covering integer programs*

Stavros Kolliopoulos, *McMaster University, Canada*

10:30 - 11:00 *Designing networks for selfish users is hard*

Tim Roughgarden, *Cornell University, USA*

11:00 - 11:30 *Classifying customer-provider relationships in the internet*

Thomas Erlebach, *ETH Zürich, Switzerland*

11:30 - 12:00 *Widesense nonblocking WDM crossconnects*

Gordon Wilfong, *Bell Labs, USA*

Monday Afternoon Session Chair: Roberto Solis-Oba

15:15 - 15:45 *Spanning Trees with a bounded number of branch vertices*

Luisa Gargano, *University of Salerno, Italy*

16:30 - 17:00 *Approximation algorithms for K-colourability*

Ingo Schiermeyer, *TU Freiberg, Germany*

17:00 - 17:30 *On approximation algorithms for data mining applications*

Foto Afrati, *TU Athens, Greece*

17:30 - 18:00 *Approximating the geometric minimum diameter spanning tree*

Alexander Wolff, *Universität Greifswald, Germany*

Tuesday Morning Session Chair: Evripidis Bampis

9:30 - 10:00 *Huffman coding with unequal letter costs*

Claire Kenyon, *Université Paris Sud, France*

10:00 - 10:30 *3 - Approximation algorithm for the k - median problem*

Roberto Solis-Oba, *University of Western Ontario, Canada*

11:00 - 11:30 *Fast approximation of minimum multicast congestion*

Anand Srivastav, *Universität Kiel, Germany*

11:30 - 12:00 *Approximation algorithms for general packing problems with logarithmic potential function*

Klaus Jansen, *Universität Kiel, Germany*

Tuesday Afternoon Session Chair: Pino Persiano

15:00 - 15:30 *Selfish resource allocation*

Paul Spirakis, *CTI Patras, Greece*

16:00 - 16:30 *Bicriteria scheduling problems with communication delays*

Evripidis Bampis, *University of Evry, France*

16:30 - 17:00 *On the approximation of a geometric set covering problem*

Frits C.R. Spieksma, *Katholieke Universiteit Leuven, Belgium*

17:00 - 17:30 *Scheduling in switching networks with set-up delays*

Ioannis Milis, *Athens University of Economic and Business, Greece*

Wednesday Morning Session Chair: Gordon Wilfong

9:30 - 10:00 *On randomized online scheduling*

Susanne Albers, *Universität Freiburg, Germany*

10:00 - 10:30 *There are spiders in dense graphs*

Mikael Hammar, *University of Salerno, Italy*

11:00 - 11:30 *Dynamic routing on networks with fixed-size buffers*

Adi Rosén, *Technion-Haifa, Israel*

11:30 - 12:00 *On-line randomized call control in cellular networks*

Ioannis Caragiannis, *University of Patras, Greece*

Thursday Morning Session Chair: Luisa Gargano

9:30 - 10:30 *Peer-to-peer networks: graphs, theorems and experiments*

Prabhakar Raghavan, *Verity Inc.-Sunnyvale, USA*

11:00 - 11:30 *TCP/IP-like protocol maximizing network-wide throughput*

Neal Young, *Akamai Technologies, USA*

11:30 - 12:00 *Approximate trade-off for bicriteria batching and parallel machine scheduling problems*

Eric Angel, *University of Evry, France*

Thursday Afternoon Session Chair: Frits Spieksma

15:00 - 15:30 *Selfish Routing: Knapsack-Like Equilibria*

Berthold Vöcking, *MPI-Saarbrücken, Germany*

16:00 - 16:30 *Online distance constrained labeling of disk graphs*

Aleksei Fishkin, *University of Kiel, Germany*

16:30 - 17:00 *Efficient communication in directed ad-hoc radio networks*

Leszek Gasieniec, *University of Liverpool, UK*

17:00 - 17:30 *Optical routing and wavelength assignment in multibuffer WDM networks*

Hervé Rivano, *INRIA-Sophia Antipolis, France*

Friday Morning Session Chair: Thomas Erlebach

09:00 - 09:30 *Finding minimum hidden guard sets*

Stephan Eidenbenz, *ETH Zürich, Switzerland*

09:30 - 10:00 *Coping with limited bandwidth in all-optical networks*

Aristeidis Pagourtzis, *ETH Zürich, Switzerland*

10:30 - 11:00 *Approximation algorithms for multicommodity flows over time*

Martin Skutella, *TU Berlin, Germany*

11:00 - 11:30 *Parallel scheduling problems in next generation - wireless networks*

Luca Becchetti, *Università di Roma, Italy*

11:30 - 12:00 *An improved communication - randomness tradeoff*

Martin Fürer, *Pennsylvania State University, USA*

3 Abstracts

Approximation Algorithms for Data Mining Applications

Foto N. Afrati
National Technical University of Athens, Greece
`afrati@softlab.ece.ntua.gr`

An overview was presented of methods and techniques used in recent research to process data stored in secondary or tertiary memory and are read into main memory only once—a.k.a data streams. Data mining problems seek to process such data. Some of the main issues in data mining research are discussed.

This talk is based on a chapter of a book to be edited by E. Bampis, K. Jansen, C. Kenyon. It will soon be in my web page:

<http://www.softlab.ntua.gr/facilities/public/AD/foto/afrati.html>

On Randomized Online Scheduling

Susanne Albers
Freiburg University, Germany
`salbers@informatik.uni-freiburg.de`

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 performs better than known deterministic algorithms for general m . Our algorithm is 1.916-competitive. We show that this performance cannot be proven for a deterministic online algorithm based on analysis techniques that have been used in the literature so far.

On the approximate tradeoff for bicriteria batching and parallel machine scheduling problems

Eric Angel
Université d'Évry, France
`angel@lami.univ-evry.fr`

We consider multiobjective scheduling problems, i.e. scheduling problems that are evaluated with respect to many cost criteria, and we are interested in determining a trade-off (*Pareto curve*) among these criteria. We study two types of bicriteria scheduling problems: *single-machine batching* problems and *parallel machine* scheduling problems. Instead of proceeding in a problem-by-problem basis, we identify a class of multiobjective optimization problems possessing a fully polynomial time approximation scheme (FPTAS) for computing an ε -approximate Pareto curve. This class contains a set of problems whose Pareto curve can be computed via a simple pseudo-polynomial dynamic program for which the objective and transition functions satisfy some, easy to verify, arithmetical conditions. Our study is based on a recent work of Woeginger [SODA'99] for the single criteria optimization ex-benevolent problems. We show how our general result can be applied to the considered scheduling problems. (Joint work with: Evripidis Bampis and Alexander Kononov)

Approximation algorithms for the bicriteria UET-UCT scheduling problem with communications delays

Evripidis Bampis
 Université d'Évry, France
 bampis@lami.univ-evry.fr

We study the problem of simultaneously minimizing the makespan and the average weighted completion time for the precedence multiprocessor constrained scheduling problem with unit execution times and unit communications delays, known as the UET-UCT problem. We propose a simple $(16/9, 16/9)$ -approximation algorithm for the problem with an unbounded number of machines. We improve our algorithm by adapting a technique first introduced by Aslam et al. and provide a $(1.745, 1.745)$ -approximate solution. The main contribution of this paper is to provide further evidence that the general techniques introduced in the paper of Stein and Wein, mainly for providing bounds on the existence of good bicriteria algorithms, may be combined with linear programming and may provide a new way of designing efficient bicriteria approximation algorithms. For the considered scheduling problem we prove the existence of a $(1.445, 1.445)$ -approximate solution for our problem, improving the generic existence result of Aslam et al. Notice also that our results for the unrestricted number of processors case hold for the more general scheduling problem with small communication delays, and for two other classical optimality criteria, maximum lateness and weighted lateness. Finally, we propose an approximation algorithm for the UET-UCT problem where the number of processors is bounded. (Joint work with Alexander Kononov)

Parallel scheduling problems in next generation - wireless networks

Luca Becchetti
Università di Roma, Italy
becchett@dis.uniroma1.it

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.

On-line Randomized Call Control in Cellular Networks

Ioannis Caragiannis
University of Patras
caragian@cti.gr

We address an important communication issue arising in cellular (wireless) networks that utilize Frequency Division Multiplexing (FDM) technology. In such networks, many users within the same geographical region (cell) can communicate simultaneously with other users of the network using distinct frequencies. The spectrum of the available frequencies is limited; thus, efficient solutions to the call control problem are essential. The objective of the call control problem is, given a spectrum of available frequencies and users that wish to communicate, to maximize the number of users that communicate without signal interference.

We study the on-line version of the problem using competitive analysis. We present lower bounds on the competitive ratio of deterministic on-line algorithms

and the main ideas of the analysis of algorithm p -Random, an intuitive randomized on-line call control algorithm. We demonstrate how the analysis techniques can be used for proving efficient competitive ratios (i.e., better than the deterministic lower bounds) on cellular networks where each cell is a regular hexagon and cellular networks which are sparse and irregular. We also present lower bounds on the competitiveness of randomized on-line call control algorithms in some typical network topologies.

The talk was based on joint work with Christos Kaklamanis and Evi Papaioannou.

Finding minimum hidden guard sets

Stephan Eidenbenz
ETH Zürich, Switzerland
eidenben@inf.ethz.ch

We study the problem MINIMUM HIDDEN GUARD SET, which consists of positioning a minimum number of guards in a given polygon or terrain such that no two guards see each other and such that every point in the polygon or on the terrain is visible from at least one guard. By constructing a gap-creating reduction from MAXIMUM 5-OCCURRENCE-3-SATISFIABILITY, we show that this problem cannot be approximated by a polynomial-time algorithm with an approximation ratio of $n^{1-\epsilon}$ for any $\epsilon > 0$, unless $NP = P$, where n is the number of polygon or terrain vertices. The result even holds for input polygons without holes. This separates the problem from other visibility problems such as guarding and hiding, where strong inapproximability results only hold for polygons with holes. Furthermore, we show that an approximation algorithm achieves a matching approximation ratio of n .

Classifying Customer-Provider Relationships in the Internet

Thomas Erlebach
ETH Zürich, Switzerland
erlebach@tik.ee.ethz.ch

The problem of inferring customer-provider relationships in the autonomous system topology of the Internet leads to the following optimization problem: given an undirected graph G and a set P of paths in G , orient the edges of G such that

as many paths as possible are valid, meaning that they do not contain an internal node with both incident edges on the path directed away from that node. The complexity of this problem was left open by Subramanian et al. (“Characterizing the Internet hierarchy from multiple vantage points,” INFOCOM 2002). We show that finding an orientation that makes all paths valid (if such an orientation exists) can be done in linear time and that the maximization version of the problem is \mathcal{NP} -hard and cannot be approximated within $1/n^{1-\varepsilon}$ for n paths unless $\mathcal{NP} = \text{co-}\mathcal{RP}$. We present constant-factor approximation algorithms for the case where the paths have bounded length; the problem is shown to be APX-hard in this case. Experimental results demonstrate that our approximation algorithm yields very good solutions on real data sets.

(This is joint work with Alexander Hall and Thomas Schank.)

Online distance constrained labeling of disk graphs

Aleksei Fishkin

Universität Kiel, Germany

`avf@informatik.uni-kiel.de`

A disk graph is the intersection graph of a set of disks in the plane. We consider the problem of assigning labels to vertices of a disk graph satisfying a sequence of distance constraints. Our objective is to minimize the distance between the least and the largest labels. We propose an on-line labeling algorithm on disk graphs, if the maximum and minimum diameters are bounded. We give the upper and lower bounds on its competitive ratio, and show that the algorithm is asymptotically optimal. In more detail we explore the case of distance constraints $(2, 1)$, and present two off-line approximation algorithms. The last one we call robust, i.e. it does not require the disks representation and either outputs a feasible labeling, or answers the input is not a unit disk graph.

An Improved Communication-Randomness Tradeoff

Martin Fürer

Pennsylvania State University, USA and ETH Zürich, Switzerland

`furerc@cse.psu.edu`

Two processors receive inputs x and y respectively. The communication complexity of the function f is the number of bits (as a function of the input size) that the processors have to exchange to compute $f(x, y)$ for worst case inputs x and y . The list-nondisjointness problem ($x = (x_1, \dots, x_n)$, $y = (y_1, \dots, y_n)$ $x_i, y_i \in \mathbf{Z}_2^n$, to decide whether $\exists i x_i = y_i$) exhibits maximal discrepancy between deterministic (n^2+1) and Las Vegas $(\Theta(n))$ communication complexity. Fleischer, Jung, Mehlhorn (1995) have shown that if a Las Vegas algorithm expects to communicate $\Omega(n \log n)$ bits, then this can be done with a small number of coin tosses.

Even with an improved (optimal) randomness efficiency, the result extends now to the (more interesting) communication efficient algorithms. R coin tosses are sufficient for $O(n/2^R)$ transmitted bits.

Spanning trees with bounded number of branch vertices

Luisa Gargano
Università di Salerno, Italy
lg@dia.unisa.it

We introduce the following combinatorial optimization problem: Given a connected graph G , find a spanning tree T of G with the smallest number of branch vertices (vertices of degree 3 or more in T). The problem is motivated by new technologies in the realm of optical networks. We investigate algorithmic and combinatorial aspects of the problem.

(Joint work with Pavol Hell, Simon Fraser University, Ladislav Stacho, Simon Fraser University and Ugo Vaccaro, Università di Salerno)

Efficient Communication in Ad-Hoc Radio Networks Leszek Gasieniec

University of Liverpool, UK
leszek@csc.liv.ac.uk

We consider the problem of efficient communication in ad-hoc radio networks. We are especially interested in ‘broadcasting’ (one-to-all communication) and ‘gossiping’ (all-to-all communication). Previous work in the field has been focused on distributed randomized broadcasting algorithms designed for unknown ad-hoc networks, and on deterministic off-line broadcasting algorithms assuming full (or at

least partial) knowledge of the radio network topology. We start with presentation of adaptive broadcasting algorithms in ad-hoc radio networks that are simultaneously distributed and deterministic. Then we show how to perform efficient deterministic and randomized radio gossiping in model with (arbitrarily) large and then small messages. We conclude the presentation with a list of open problems in the field.

There are spiders in dense graphs

Mikael Hammar
Università di Salerno, Italy
hammar@dia.unisa.it

Given a graph in which the degree sum of any $k + 2$ independent vertices is at least $n - 1$, does it contain a spanning tree with at most k vertices of degree higher than two? This has recently been conjectured by Gargano et al., and we show the conjecture for the special case when $k = 1$. A tree with only one vertex of degree higher than two is called a spider, and so we can guarantee the existence of spanning spiders in dense graphs. For the bipartite case we show that there is a spanning spider in $G = (U, V, E)$ if $|U| \leq |V|$, and if for any pair u, v with $u \in U$ and $v \in V$, $d(v) \geq (|U| + 1)/2$, and $d(u) + d(v) \geq |V| + 1$. In addition we show that the bound in the bipartite case is the best possible.

Approximation Algorithms for General Packing Problems with Logarithmic Potential Function

Klaus Jansen
Universität Kiel, Germany
kj@informatik.uni-kiel.de

In this talk we present an approximation algorithm based on Lagrangian decomposition via a logarithmic potential reduction to solve a general packing or min-max resource sharing problem with M nonnegative convex constraints on a convex set B . We generalize a method by Grigoriadis and Khachiyan to the case with weak approximate block solvers (i.e. with only constant, logarithmic or even worse approximation ratios). We show that the algorithm needs at most $O(M(\epsilon^{-2} \ln \epsilon^{-1} + \ln M))$ calls to the block solver, a bound independent on the data and the approximation ratio of the block solver. For small approximation ratios the algorithm needs at

most $O(M(\epsilon^{-2} + \ln M))$ calls to the block solver. As an application we study the problem of minimizing the maximum edge congestion in a multicast communication network. Interestingly the block problem here is the classical minimum Steiner tree problem that can be solved only approximately. We show how to use approximation algorithms for the minimum Steiner tree problem to solve the edge congestion problem approximately.

This is joint work with Hu Zhang, University of Kiel.

Huffman Coding with Unequal Letter Costs

Claire Kenyon
 Université Paris Sud
 enyon@lri.fr

In the standard Huffman coding problem, one is given a set of words and for each word a positive frequency. The goal is to encode each word w as a codeword $c(w)$ over a given alphabet. The encoding must be prefix free (no codeword is a prefix of any other) and should minimize the weighted average codeword size $\sum_w \text{freq}(w) \text{size}(c(w))$. The problem has a well-known polynomial-time algorithm due to Huffman.

Here we consider the generalization in which the letters of the encoding alphabet may have non-uniform lengths, as in the dots and dashes of Morse Code. The goal is to minimize the weighted average codeword length $\sum_w \text{freq}(w) \text{cost}(c(w))$, where $\text{cost}(s)$ is the sum of the (possibly non-uniform) lengths of the letters in s . Despite much previous work, the problem is not known to be NP-hard, nor was it previously known to have a polynomial-time approximation algorithm. Here we describe a polynomial-time approximation scheme (PTAS) for the problem.

(Joint work with Mordecai Golin and Neal Young.)

Tight Approximation Results for General Covering Integer Programs

Stavros G. Kolliopoulos
 McMaster University
 stavros@mcmaster.ca

We study approximation algorithms for solving a general *covering integer program*. An n -vector x of nonnegative integers is sought, which minimizes $c^T \cdot x$,

subject to $Ax \geq b$, $x \leq d$. The entries of A, b, c are nonnegative. Let m be the number of rows of A . Covering problems have been heavily studied in combinatorial optimization. We focus on the effect of the *multiplicity constraints*, $x \leq d$, on approximability. Two longstanding open questions remain for this general formulation with upper bounds on the variables.

- (i) The integrality gap of the standard LP relaxation is arbitrarily large. Existing approximation algorithms that achieve the well-known $O(\log m)$ -approximation with respect to the LP value do so at the expense of violating the upper bounds on the variables by the same $O(\log m)$ multiplicative factor. What is the smallest possible violation of the upper bounds that still achieves cost within $O(\log m)$ of the standard LP optimum?
- (ii) The best known approximation ratio for the problem has been $O(\log(\max_j \sum_i A_{ij}))$ since 1982. This bound can be as bad as polynomial in the input size. Is an $O(\log m)$ -approximation, like the one known for the special case of Set Cover, possible?

We settle these two open questions. To answer the first question we give an algorithm based on the relatively simple new idea of randomly rounding variables to smaller-than-integer units. To settle the second question we give a reduction from approximating the problem while respecting multiplicity constraints to approximating the problem with a bounded violation of the multiplicity constraints.

(Joint work with Neal Young.)

Scheduling in switching networks with set-up delays

Ioannis Milis

Athens University of Economics and Business, Greece

milis@aueb.gr

The *preemptive bipartite scheduling problem (PBS)* arises in switching communication systems, where each input and output port can be involved in at most one communication at the same time. Given a set of communication tasks to be transmitted from the input to the output ports of such a system, we aim to find a schedule minimizing the overall transmission time. To achieve this, the preemption of communication tasks is allowed. However, in practice preemption comes with a set-up delay, d , and this renders the problem NP-hard.

We present a $2 - \frac{1}{d+1}$ approximation algorithm for this problem and we also show that preempting every α time units, can not lead to a better approximation ratio for any α .

(Joint work with Foto Afrati, Timos Aslanidis and Evripidis Bampis.)

Coping with Limited Bandwidth in All-Optical Networks

Aris Pagourtzis
ETH Zürich, Switzerland
pagour@inf.ethz.ch

Limited bandwidth in WDM all-optical networks calls for algorithms that may help to maximize the number of established connections. We discuss the following problem: given a set of requests over a network and a number of available wavelengths, find a maximum number of requests that can be assigned a path and a color so that edge-intersecting paths have different colors.

We propose approximation algorithms for ring networks under both the undirected and the directed settings, corresponding to symmetric and one-way communication respectively. The new algorithms guarantee approximation ratio $2/3$ for the undirected case and $7/11$ for the directed case, thus improving upon the $(e-1)/e$ approximation ratio which was the best known so far for both cases.

Along the way we introduce the w-Balanced Matching problem and establish the following: a r -approximation algorithm for w-Balanced Matching yields a $(r+1)/(r+2)$ -approximation algorithm for the problem of maximizing the number of satisfied directed requests over a bidirected ring. We give a $3/4$ -approximation algorithm for this new matching problem.

(Joint work with Christos Nomikos and Stathis Zachos.)

Peer-to-peer networks: graphs, theorems and experiments

Prabhakar Raghavan
Verity Inc.
pragh@verity.com

This lecture covers two topics related to peer-to-peer networks. The first concerns a random graph construction process in which nodes arrive and depart over time,

connecting to nodes in the network chosen using only "local" knowledge. The process nevertheless ensures that with high probability the resulting network has small diameter. In the second half of the lecture we discuss "topic-segmented topologies" in which peer-to-peer networks are organized based on the contents of the nodes. The organization ensures that nodes with similar content are close to one another in the network. Consequently, queries are not broadcast through the network but rather matched to a topic and despatched to the appropriate region of the network. This results in reduced network traffic for a given level of content recall - we describe experimental results to this effect.

Lightpath assignement in multifiber WDM networks with wavelength translation

Hervé Rivano
 INRIA, France
 Herve.Rivano@sophia.inria.fr

We consider the problem of finding a lightpath assignment for a given set of communication requests on a multifiber WDM optical network with wavelength translators. Given such a network, and w the number of wavelengths available on each fiber, $k(e)$ the number of fiber on each link e and $c(u)$ the number of partial wavelength translation available on each node u , our problem stands for deciding whether it is possible to find a w -lightpath for each request in the set such that there is no link e carrying more that $k(e)$ lightpaths using the same wavelength nor node u where more than $c(u)$ wavelength translations take place. Our main theoretical result is the writing of this problem as a particular instance of integral multicommodity flow, hence integrating routing and wavelength assignment in the same model. We then provide two heuristics mainly based upon enhancement of the classical randomized rounding of fractional multicommodity flow that are two different answers to the trade-off between efficiency and tightness of approximation and discuss their practical performances on both theoretical and real-world instances.

Dynamic Routing on Networks with Fixed-Size Buffers

Adi Rosén
 Technion, Israel
 adiro@cs.technion.ac.il

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 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 greedy, local-control, on-line 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 W. Aiello, E. Kushilevitz, and R. Ostrovsky

Designing Networks for Selfish Users is Hard

Tim Roughgarden
Cornell University, USA
timr@cs.cornell.edu

Given a network with congestion-dependent edge latencies and a prescribed source-destination pair and traffic rate, which subnetwork will exhibit the best performance when used selfishly? Braess's Paradox shows that the trivial algorithm of choosing the whole network is a suboptimal heuristic; we show that it is the best possible polynomial-time heuristic, unless $P=NP$. We also give an infinite family of examples generalizing Braess's Paradox, providing the first demonstration that the severity of the paradox grows with the network size.

Forbidden Subgraphs and 3-Colourability

Ingo Schiermeyer
Technische Universität Bergakademie Freiberg, Germany
schierme@math.tu-freiberg.de

The 3-colourability problem is a well-known NP-complete problem. It remains NP-complete for triangle-free graphs of maximum degree 4 and for claw-free graphs. Sumner has shown that triangle-free and P_5 -free or triangle-free, P_6 -free and C_6 -free graphs are 3-colourable.

We present polynomial time algorithms to colour a (K_3, P_5) -free graph with three colours and a (K_3, P_6) -free graph with four colours. Furthermore we show that (after suitable reductions) every 4-chromatic (K_3, P_6) -free graph G contains the Mycielski-Grötzsch graph as an induced subgraph and is a subgraph of the Clebsch graph.

Using small dominating sets we show that 3-colourability can be decided and a corresponding 3-colouring can be determined in polynomial time for the class of P_6 -free graphs.

3-colourability can be also decided and a corresponding 3-colouring can be determined in polynomial time for the class of claw-free and hourglass-free graphs $(K_{1,3}, K_1 + 2K_2)$ and claw-free and t -spider-free graphs (a $K_{1,t}$ with each edge subdivided).

(Joint work with Bert Randerath and Meike Tewes, Technische Universität Bergakademie Freiberg.)

Approximation Algorithms for Multicommodity Flows Over Time

Martin Skutella
 TU Berlin, Germany
 skutella@math.tu-berlin.de

Traditionally, flows over time are solved in time-expanded networks which contain one copy of the original network for each discrete time step. While this method makes available the whole algorithmic toolbox developed for static flows, its main and often fatal drawback is the enormous size of the time-expanded network. In particular, this approach usually does not lead to efficient algorithms with running time polynomial in the input size since the size of the time-expanded network is only pseudo-polynomial.

We present two different approaches for coping with this difficulty. Firstly, inspired by the work of Ford and Fulkerson on maximal s - t -flows over time (or ‘maximal dynamic s - t -flows’), we show that static, length-bounded flows lead to provably good multicommodity flows over time. These solutions not only feature a simple structure but can also be computed very efficiently in polynomial time.

Secondly, we investigate ‘condensed’ time-expanded networks which rely on a rougher discretization of time. Unfortunately, there is a natural tradeoff between the roughness of the discretization and the quality of the achievable solutions. However,

we prove that a solution of arbitrary precision can be computed in polynomial time through an appropriate discretization leading to a condensed time expanded network of polynomial size. In particular, this approach yields a fully polynomial time approximation scheme for the quickest multicommodity flow problem and also for more general problems.

This is joint work with Lisa Fleischer.

An approximation algorithm for the k-median problem

Roberto Solis-Oba
University of Western Ontario, Canada
`solis@csd.uwo.ca`

The k-median problem is a central problem in Operations Research that has captured the attention of the algorithms community in recent years. Despite its importance, a non-trivial approximation algorithm for the problem eluded researchers for a long time.

Remarkably, a succession of papers with ever improved performance ratios have been designed in the last few years. In this talk we described a new approximation algorithm for the k -median problem based on the algorithm by Jain and Vazirani. Our algorithm avoids the complex rounding scheme required by the above algorithm and the algorithm of Charikar and Guha.

Approximation of a geometric set-covering problem

Frits Spieksma
Katholieke Universiteit Eichstätt, Netherlands
`frits.spieksma@unimaas.nl`

In this talk we addressed a packing and a covering problem. The packing problem can be viewed as finding a maximum stable set on a graph that is the edge-union of an interval graph and a graph that is a disjoint union of cliques. Let k be the maximum number of nodes in such a clique. The covering problem amounts to finding a minimum clique cover in this graph. Both problems can be viewed in a geometric context.

Each problem is APX-hard. We show that the covering problem can be approximated within $2k/(k+1)$. This is done by rounding the LP-relaxation.

This is joint work with Sofia Kovaleva from Maastricht University.

Selfish resource allocation

Pavlos Spirakis
University of Patras
spirakis@cti.gr

In this talk , we propose a study of the relation of classical game theory and algorithmic complexity via a case study. The case study is a simple routing problem of n agents with loads and m parallel links . This is a model of non-cooperating agents proposed originally by Papadimitriou and Koutsoupias. We start from Nash equilibria which represent logical ways of behaviour from the antagonistic agents. We prove that pure equilibria exist in this game and can be found in poly-time. On the other hand , we show that it is hard to compute the best or worst pure equilibrium unless $P=NP$.

We estimate the value of equilibria via their social cost , ie the makespan they cause by scheduling the agents loads into the links. This cost is , in general , a random variable for mixed strategies , so we use its expected value. We show that it is sharp-P hard to compute the exact value of the social cost , when the probabilities of the equilibrium are given.

We also show that there are certain equilibria (that consider all links possible with nonzero probabilities) that are very close to worst case equilibria wrt the social cost. e then give a polynomial time algorithm which delivers one such equilibrium.

This is one of the first studies in the direction of algorithmic and complexity network research motivated by game theory. The work is done in cooperation with M Mavronicolas , E Koutsoupias , D Fotakis and S. Kontogiannis.

Fast approximation of minimum multicast congestion

Anand Srivastav
Universität Kiel, Germany
asr@numerik.uni-kiel.de

The problem of minimizing the maximum edge congestion in a multicast communication network with n nodes and m edges generalizes the well-known NP -hard standard routing problem. The previously best approximation algorithm known is due to Carr and Vempala (STOC 2000), and it achieves an approximation within $O(opt + \log n)$ in $\tilde{O}(n^7)$ -time. We present a new and fast approximation algorithm with complexity of only $\tilde{O}(k^2m)$, and a slightly better approximation guarantee.

We prove this result by designing a fast combinatorial algorithm for solving an LP-relaxation of the minimum multicast congestion problem generalizing the algorithm of Klein, Plotkin, Stein, and Tardos (1994) from paths to trees.

Selfish Routing: Knapsack-Like Equilibria

Berthold Vöcking
MPI für Informatik
voecking@mpi-sb.mpg.de

What is the complexity of computing the worst knapsack packing? — To be more precise, given n objects with weights w_i and profits p_i and a capacity threshold b , we ask for the complexity of computing a maximal knapsack with minimum profit (min-max knapsack). We investigate this and other more general knapsack-like optimization problems allowing different thresholds for different objects. Our study is motivated by the complexity of best- and worst-case microeconomical equilibria describing customer allocations for a service provider.

As an example of our results, we show that there is a pseudo-polynomial time algorithm for the min-max knapsack problem that is only slightly more complicated than the well known algorithm for the maximum knapsack problem. Unlike the maximum knapsack problem, however, the min-max problem does not admit an FPTAS. In fact, we can show that this problem cannot be approximated within any polynomial time computable factor.

To obtain our inapproximability results, weights and thresholds in reductions must be chosen in a very exact way. In contrast, our micro-economical motivation yields weights and thresholds that are inherently fuzzy. The recently introduced concept of smoothed algorithm analysis seems to be exactly the right methodology to deal with this fuzziness and to escape the inapproximability dilemma. As our main result, we present an approximation scheme with polynomial time smoothed-complexity for the most general variants among the knapsack-like equilibria problems under consideration.

Joint work with Rene Beier and Piotr Krysta

Wide-sense Nonblocking WDM Cross-connects

Gordon Wilfong
Bell Labs, USA
gtw@research.bell-labs.com

We consider the problem of minimizing the number of wavelength interchangers in the design of wide-sense nonblocking cross-connects for wavelength division multiplexed (WDM) optical networks. The problem is modeled as a graph theoretic problem that we call *dynamic edge coloring*. In dynamic edge coloring the nodes of a graph are fixed but edges appear and disappear, and must be colored at the time of appearance without assigning the same color to adjacent edges.

For wide-sense nonblocking WDM cross-connects with k input and k output fibers, it is straightforward to show that $2k-1$ wavelength interchangers are always sufficient. We show that there is a constant $c > 0$ such that if there are at least ck^2 wavelengths then $2k-1$ wavelength interchangers are also necessary. This improves previous exponential bounds. When there are only 2 or 3 wavelengths available, we show that far fewer than $2k-1$ wavelength interchangers are needed. However we also prove that for any $\varepsilon > 0$ and $k > 1/2\varepsilon$, if the number of wavelengths is at least $1/\varepsilon^2$ then $2(1-\varepsilon)k$ wavelength interchangers are needed.

This is joint work with P. Haxell (U. Waterloo), A. Rasala (MIT) and P. Winkler (Bell Labs).

Approximating the Geometric Minimum-Diameter Spanning Tree

Alexander Wolff
 Universität Greifswald
 awolff@mail.uni-greifswald.de

Let P be a set of n points in the plane. The geometric minimum-diameter spanning tree (MDST) of P is a tree that spans P and minimizes the Euclidian length of the longest path. Ho et al. showed that there is always a mono- or a dipolar MDST. (A pole is a vertex of degree greater than 1.) The more difficult dipolar case can so far only be solved in $O(n^{3-1/6})$ time. This is due to Chan who improved a cubic-time algorithm of Ho et al. by using semi-dynamic data structures.

We gave some fast approximation schemes for the MDST, i.e. factor- $(1 + \epsilon)$ approximation algorithms. One uses a grid and takes $O^*(\epsilon^{-6} + n)$ time, where the O^* -notation hides terms of order $\log \epsilon^{-1}$. The other uses a geometric data structure, the well-separated pair decomposition, and takes $O(n\epsilon^{-3} + En \log n)$ time. A combination of the two approaches runs in $O^*(\epsilon^{-5} + n)$ time, i.e. it is a strongly linear-time approximation scheme (LTAS) of order 5.

Interesting questions are as follows: Is there a LTAS of order less than 5? Can the exact MDST be computed faster than by Chan's algorithm? What is a lower bound for the time-complexity of computing the exact MDST?

This research was joint work with Joachim Gudmundsson, Herman Haverkort, Sang-Min Park, and Chan-Su Shin.

On-line, end-to-end congestion control

Neal E. Young
Akamai Technologies
`neal@acm.org`

We study the convergence complexity of a simple, end-to-end, multiplicative-increase/multiplicative-decrease protocol for Internet congestion control. We show the protocol achieves near-optimal network-wide throughput (w.r.t. the static optimum) provided every connection stays open for at least logarithmically many rounds. The analysis leverages recent work on Lagrangian-relaxation algorithms for packing and covering problems. The protocol is implementable in the current Internet by modifying the TCP server. The protocol has a specific practical application – bandwidth testing, and, for quality-of-service (QOS) applications, generalizes to maximize total weighted throughput (with user-specified weights).

Joint work with Naveen Garg

Draft available at <http://arxiv.org/abs/cs.DS/0205032> .