



## ECCS Workshop: Optimisation in Complex Networks

September 28-29, 2006 / Said Business School, University of Oxford  
Organisers: Paul Spirakis (CTI Patras), Berthold Vöcking (RWTH Aachen)

Information systems like the Internet, the World Wide Web, telephone networks, mobile ad-hoc networks, or peer-to-peer networks have reached a level that puts them beyond our ability to deploy, manage, and keep them functioning correctly through traditional techniques. Reasons for this are their sheer size with millions of users and interconnected devices and their dynamics; they evolve dynamically over time, i.e., components change or are removed or inserted permanently. Most of the existing and foreseen complex networks are furthermore built, operated and used by a multitude of diverse economic interests. For such systems, we often have to abandon the goal of global optimality and instead make use of approximation algorithms that can guarantee only to find a solution that is close to optimal or heuristics that guarantee optimality only from a local perspective, e.g., in form of a Nash equilibrium. The goal of this workshop is to present and discuss different optimisation methods and techniques that are able to cope with the possibly conflicting objectives of the participating entities.

The focus of the workshop lies on the following topics:

- scheduling and resource allocation problems
- decision making with multiple objectives
- distributed algorithms for wireless/sensor networks
- game theoretic approaches and selfish optimisation
- approximation and randomised algorithms
- mathematical analysis of heuristics

## Schedule

The workshop is scheduled on Thursday, September 28th, from 9:00 am to 6:00 pm and on Friday, September 29th, from 9:00 am to 12:30 pm in the *James Martin Institute Seminar Room*. The program with 15 talks is centred around the following invited presentations:

Artur Czumaj, University of Warwick  
Sublinear-Time Algorithms  
Thursday, 10:15 am - 11:00 am

Martin Skutella, Dortmund University  
Solving Evacuation Problems Efficiently: Earliest Arrival Flows with Multiple Sources  
Thursday, 3:15 pm - 4:00 pm

Carsten Witt, Dortmund University  
Runtime Analysis of a Simple Ant Colony Optimization Algorithm  
Friday, 10:15 am - 11:00 am

Time	September 28th	September 29th
9:00 - 9:30	M. Kutylowski	H. Ackermann
9:30 - 10:00	K. Tiemann	V. Papadopoulou
10:15 - 11:00	A. Czumaj	C. Witt
11:00 - 11:30	<i>coffee break</i>	<i>coffee break</i>
11:30 - 12:00	M. J. Serna	F. Schoppmann
12:00 - 12:30	S. Kontogiannis	P. Spirakis
12:30 - 14:00	<i>lunch</i>	<i>lunch</i>
14:00 - 14:30	A. Kaporis	-
14:30 - 15:00	Y. Bleischwitz	-
15:15 - 16:00	M. Skutella	-
16:00 - 16:30	<i>afternoon tea</i>	-
16:30 - 17:00	B. Vöcking	-
17:00 - 17:30	M. Klonowski	-

## List of Participants

1. Heiner Ackermann, RWTH Aachen
2. Yvonne Bleischwitz, University of Paderborn
3. Artur Czumaj, University of Warwick
4. Alexios Kaporis, CTI Patras
5. Marek Klonowski, TU Wroclaw
6. Spyros Kontogiannis, CTI Patras
7. Jaroslaw Kutylowski, University of Paderborn
8. Mirosław Kutylowski, TU Wroclaw
9. Matus Mihalak, University of Leicester
10. Vicky Papadopoulou, University of Cyprus
11. Maria Jose Serna, University of Catalunya, Barcelona
12. Martin Skutella, University Dortmund
13. Florian Schoppmann, University of Paderborn
14. Paul Spirakis, CTI Patras
15. Karsten Tiemann, University of Paderborn
16. Berthold Vöcking, RWTH Aachen
17. Carsten Witt, University Dortmund

## **Scheduling Communication in Ad Hoc Networks**

Mirosław Kutylowski, TU Wrocław

We consider problem of allocation of communication channel to units of an ad hoc network that have to send data. We propose slight changes in the protocols used so far and reduce significantly communication overhead due to coordination of the senders. Apart from experimental data, we develop tools for mathematical analysis of performance of these algorithms.

Joint work with Jacek Cichon, Marcin Zawada

## **Routing Splittable Flow in Games with Player-Specific Linear Latency Functions**

Carsten Tiemann, University of Paderborn

We study weighted network routing games with player-specific latency functions where  $n$  selfish players wish to route their traffic through a shared network. Most of our results are obtained for linear latency functions. We consider the case of splittable traffic. We derive several results on the existence and computational complexity of equilibria and on the price of anarchy. Our main findings are as follows:

1. For routing games on parallel links with linear latency functions without a constant term we introduce a new potential function. We show that our convex potential function is minimized if and only if the corresponding assignment is an equilibrium. This result implies that an equilibrium can be computed in polynomial time. We also show for several generalizations of the above games that such potential functions do not exist.
2. We prove upper and lower bounds on the price of anarchy for games with linear latency functions.

## **Sublinear-time algorithms**

Artur Czumaj, University of Warwick

The design of algorithms operating on massive data sets, has received a lot of attention in recent years. The practical motivation of this study is that polynomial algorithms that are efficient in relatively small inputs, may become impractical for input sizes of several gigabytes. Managing and analyzing such data sets forces us to revisit the traditional notions of efficient algorithms. For example, when we consider approximation algorithms for clustering problems in metric spaces with  $n$  points, then since their input size is  $\Theta(n^2)$ , they typically have  $\Omega(n^2)$  running time. Clearly, such a running time is not feasible for massive data sets. Constructing a sublinear time algorithm may seem to be an impossible task since it allows one to read only a small fraction of the input. However, in recent years, we have seen development of sublinear time algorithms for optimization problems arising in such diverse areas as graph theory, geometry, algebraic computations, and computer graphics.

In this talk, we will present a few examples of sublinear algorithms recently developed. Our main focus will be on techniques used to design sublinear-time algorithms to estimate basic graph parameters (e.g., the average degree of the graph or the cost of its minimum spanning tree) and for clustering problems.

## **Problems on Games: Complexity versus Succinctness**

Maria Jose Serna, University of Catalunya, Barcelona

We are interested in the study of the computational complexity of problems involving games with a large number of player's or actions. In particular in deciding the existence of a Pure Nash Equilibrium or deciding whether two strategic games are isomorphic. For doing so it is fundamental to decide how to represent a game as the input to a computer. We propose three ways of representing a game, each of them with different degree of succinctness for the components of the game and show how the computational complexity of the problems changes depending on the representation.

Joint work with Carme Àlvarez, Joaquim Gabarró, and Alina García

## Atomic Congestion Games Among Coalitions

Spyros Kontogiannis, CTI Patras

We consider algorithmic questions concerning the existence, tractability and quality of atomic congestion games, among users that are considered to participate in (static) selfish coalitions. We carefully define a coalitional congestion model among atomic players.

Our findings in this model are quite interesting, in the sense that we demonstrate many similarities with the non-cooperative case. For example, there exist potentials proving the existence of Pure Nash Equilibria (PNE) in the (even unrelated) parallel links setting; the Finite Improvement Property collapses as soon as we depart from linear delays, but there is an exact potential (and thus PNE) for the case of linear delays, in the network setting; the Price of Anarchy on identical parallel links demonstrates a quite surprising threshold behavior: it persists on being asymptotically equal to that in the case of the non-cooperative KP-model, unless we enforce a sublogarithmic number of coalitions.

But we also show crucial differences, mainly concerning the hardness of algorithmic problems that are solved efficiently in the non-cooperative case. Although we demonstrate convergence to robust PNE, we also prove the hardness of computing them. On the other hand, we can easily construct a generalized fully mixed Nash Equilibrium. Finally, we propose a new improvement policy that converges to PNE that are robust against (even dynamically forming) coalitions of small size, in pseudo-polynomial time.

Joint work with D. Fotakis, S. Kontogiannis, P. Spirakis.

## Typical case optimization problems on large graphs

Alexios Kaporis, CTI Patras

The Erdős-Renyi models of random graphs have attracted a lot of interest lately for the typical case analysis of algorithms. Proving that a problem is NP-hard or SNP-hard ( $C$ -inapproximable) usually relies on some "pathological graphs". A nice feature, shared by certain kinds of random graphs, is that almost always are not the pathological ones. Therefore, their typical behavior is, by far, non "NP-pessimistic".

We study how the typical solution behaves when compared to the optimal one for some well known optimization problems. The techniques presented have both algorithmic and combinatorial perspective.

## **Cost Sharing for Scheduling Problems - Results and Approaches**

Yvonne Bleischwitz, University of Paderborn

We consider the problem of sharing the cost of scheduling jobs on parallel machines among a set of agents. In our setting, each agent owns one job and the cost is given by the makespan of the computed assignment. We focus on group-strategyproof and budget-balanced mechanisms that do not allow free riders. Considering cross-monotonic cost-sharing methods that imply group-strategyproofness, we can only obtain approximative budget balance. It is an open problem, if group-strategyproofness can be achieved without these cross-monotonic cost-sharing methods for certain kinds of scheduling problems. We give a group-strategyproof mechanism for a subproblem and point to difficulties in designing such mechanisms for general scheduling problems.

## **Solving Evacuation Problems Efficiently: Earliest Arrival Flows with Multiple Sources**

Martin Skutella, University Dortmund

Earliest arrival flows capture the essence of evacuation planning. Given a network with capacities and transit times on the arcs, a subset of source nodes with supplies and a sink node, the task is to send the given supplies from the sources to the sink "as quickly as possible". The latter requirement is made more precise by the earliest arrival property which requires that the total amount of flow that has arrived at the sink is maximal for all points in time simultaneously. It is a classical result from the 1970s that, for the special case of a single source node, earliest arrival flows do exist and can be computed by essentially applying the Successive Shortest Path Algorithm for min-cost flow computations. While it has previously been observed that an earliest arrival flow still exists for multiple sources, the problem of computing one efficiently has been open for many years. We present an exact algorithm for this problem whose running time is strongly polynomial in the input plus output size of the problem.

# Worst Case and Probabilistic Analysis of the 2-Opt Algorithm for the TSP

Berthold Vöcking, RWTH Aachen

*2-Opt* is probably the most basic and widely used local search heuristic for the TSP. This heuristic achieves amazingly good results on “real world” Euclidean instances both with respect to running time and approximation ratio. There are numerous experimental studies on the performance of *2-Opt*. However, the theoretical knowledge about this heuristic is still very limited. Not even its worst case running time on Euclidean instances was known so far. In this paper, we clarify this issue by presenting a family of Euclidean instances on which *2-Opt* can take an exponential number of steps.

Previous probabilistic analyses were restricted to instances in which  $n$  points are placed uniformly at random in the unit square  $[0, 1]^2$ , where it was shown that the expected number of steps is bounded by  $\tilde{O}(n^{10})$  for Euclidean instances. We consider a more advanced model of probabilistic instances in which the points can be placed according to general distributions on  $[0, 1]^2$ . In particular, we allow different distributions for different points. We study the expected running time in terms of the number  $n$  of points and the maximal density  $\phi$  of the probability distributions. We show an upper bound on the expected length of any *2-Opt* improvement path of  $\tilde{O}(n^{4+1/3} \cdot \phi^{8/3})$ . When starting with an initial tour computed by an insertion heuristic, the upper bound on the expected number of steps improves even to  $\tilde{O}(n^{3+5/6} \cdot \phi^{8/3})$ . If the distances are measured according to the Manhattan metric, then the expected number of steps is bounded by  $\tilde{O}(n^{3.5} \cdot \phi)$ . In addition, we prove an upper bound of  $O(\sqrt{\phi})$  on the expected approximation factor w. r. t. both of these metrics.

Let us remark that our probabilistic analysis covers as special cases the uniform input model with  $\phi = 1$  and a smoothed analysis with Gaussian perturbations of standard deviation  $\sigma$  with  $\phi \sim 1/\sigma^2$ . Besides random metric instances, we also consider an alternative random input model in which an adversary specifies a graph and distributions for the edge lengths in this graph. In this model, we achieve even better results on the expected running time of *2-Opt*.

## Number of Nodes in Chord Protocol

Marek Klonowski, TU Wroclaw

We present some remarks about widely-known Chord protocol. We prove that it is strongly unbalanced structure - i.e. there are significant number of nodes that has very small or very large load in comparison with the average. We also show two fairly practical algorithms that can be used for estimating number of nodes in the network. All results presented are provided with rigid mathematical proofs. However, some described phenomena here are supported by experimental results.

Joint work with Jacek Cichon, Marek Klonowski, Przemyslaw Kobylanki Lukasz Krzywiecki, Bartomiej Bartlomiej Rozanski and Pawel Zielinski.

## On the Impact of Combinatorial Structure on Congestion Games and their Extensions

Heiner Ackermann, RWTH Aachen

We study the impact of combinatorial structure on basic properties of congestion games and their extensions. To be precise, we consider the following two questions.

1. Is there a necessary and sufficient condition on the players' strategy spaces that guarantees polynomial time convergence in congestion games?
2. Is there a necessary and sufficient condition on the players' strategy spaces that guarantees the existence of pure Nash equilibria in weighted and in player-specific congestion games?

We show that in both cases the *matroid property*, i. e., the strategy space of each player consists of the set of bases over the set of resources, is the required condition. To show sufficiency, we generalize results on singleton congestion games, i. e. the strategy space of each player consists of singleton sets only, towards matroid congestion games. To show necessity, we present a characterization of non-matroid strategy spaces that allows to construct counterexamples in both questions.

## Games with Attackers and a Defender

Vicky Papadopoulou, University of Cyprus

We survey a research line recently initiated by Mavronicolas *et al.* [ISAAC 2005], concerning a *strategic game* on a graph  $G(V, E)$  with two confronting classes of randomized players:  $\nu$  *attackers* who choose vertices and wish to minimize the probability of being caught by the *defender*, who chooses edges and gains the expected number of attackers it catches. So, the defender captures system rationality. In a *Nash equilibrium*, no single player has an incentive to unilaterally deviate from its randomized strategy. The *Price of Defense* is the worst-case ratio, over all Nash equilibria, of the *optimal* gain of the defender (which is  $\nu$ ) over the gain of the defender at a Nash equilibrium. We present a comprehensive collection of trade-offs between the Price of Defense and the computational efficiency of Nash equilibria.

- We present an algebraic characterization of (*mixed*) Nash equilibria.
- No (non-trivial) instance of the graph-theoretic game has a *pure* Nash equilibrium. This is an immediate consequence of some *covering* properties proved for the *supports* of the players in all (*mixed*) Nash equilibria.
- We present a reduction of the game to a *Zero-Sum Two-Players Game* that proves that a general Nash equilibrium can be computed via Linear Programming in polynomial time. However, the reduction does not provide any apparent guarantees on the Price of Defense.
- To obtain guarantees on Price of Defense, we present an analysis of several structured Nash equilibria:
- In a *Matching Nash equilibrium*, the support of the defender is an *Edge Cover* of the graph. Matching Nash equilibria are shown to still be computable in polynomial time, and that they incur a Price of Defense of  $\alpha(G)$ , the *Independence Number* of  $G$ .
- In a *Perfect Matching Nash equilibrium*, the support of the defender is a *Perfect Matching* of the graph. Perfect Matching Nash equilibria are shown to be computable in polynomial time, and that they incur a Price of Defense of  $\frac{|V|}{2}$ .
- We consider a generalization of the basic model with an increased power for the defender: it is able to scan a simple *path* of the network instead of a single edge. Deciding existence of a pure Nash equilibrium is shown to be an NP-complete problem for this model.

Joint work with Marios Mavronicolas, Anna Philippou, and Paul Spirakis

## **Runtime Analysis of a Simple Ant Colony Optimization Algorithm**

Carsten Witt, University Dortmund

Ant Colony Optimization (ACO) has become quite popular in recent years. In contrast to many successful applications, the theoretical foundation of this randomized search heuristic is rather weak. Building up such a theory is demanded to understand how these heuristics work as well as to come up with better algorithms for certain problems. Up to now, only convergence results have been achieved showing that optimal solutions can be obtained in finite time. We present the first runtime analysis of an ACO algorithm, which transfers many rigorous results on the runtime of a simple evolutionary algorithm to our algorithm. Moreover, we examine the choice of the evaporation factor, a crucial parameter in ACO algorithms, in detail for a toy problem. By deriving new lower bounds on the tails of sums of independent Poisson trials, we determine the effect of the evaporation factor almost completely and prove a phase transition from exponential to polynomial runtime.

Joint work with Frank Neumann.

## **Wardrop Equilibria and Price of Stability for Bottleneck Games with Splittable Traffic**

Florian Schoppmann, University of Paderborn

We look at the scenario of having to route a continuous rate of traffic from a source node to a sink node in a network, where the objective is to maximize throughput. This is of interest, e.g., for providers of streaming content in communication networks. The overall path latency, which was relevant in other non-cooperative network routing games such as the classic Wardrop model, is of lesser concern here.

To that end, we define bottleneck games with splittable traffic where the throughput on a path is inversely proportional to the maximum latency of an edge on that very path—the bottleneck latency. Therefore, we define a Wardrop equilibrium as a traffic distribution where this bottleneck latency is at minimum on all used paths. As a measure for the overall system well-being—called social cost—we take the weighted sum of the bottleneck latencies of all paths.

Our main findings are as follows: First, we prove social cost of Wardrop equilibria on series parallel graphs to be unique. Even more, for any graph whose subgraph induced by all simple start-destination paths is not series parallel, there exist games having equilibria with different social cost. For the price of stability, we give an independence result with regard to the network topology. Finally, our main result is giving a new exact price of stability for Wardrop/bottleneck games on parallel links with  $M/M/1$  latency functions. This result is at the same time the exact price of stability for bottleneck games on general graphs.

## The price of Optimum of Stackelberg games on arbitrary nets and arbitrary latency functions

Paul Spirakis, CTI Patras

Let  $M$  be a single  $s$ - $t$  network of parallel links with load dependent latency functions shared by an infinite number of selfish users. This may yield a Nash equilibrium with unbounded *Coordination Ratio*. A *Leader* can decrease the coordination ratio by assigning flow  $\alpha \cdot r$  on  $M$ , and then all *Followers* assign selfishly the  $(1 - \alpha) \cdot r$  remaining flow. This is a Stackelberg Scheduling Instance  $(M, r, \alpha)$ ,  $0 \leq \alpha \leq 1$ . It was shown that it is weakly NP-hard to compute the optimal Leaders strategy.

For any such network  $M$  we efficiently compute the minimum portion  $\beta \cdot M$  of flow  $r$  needed by a Leader to induce  $M$ 's optimum cost, as well as his optimal strategy.

Unfortunately, Stackelberg routing in more general nets can be arbitrarily hard. Roughgarden presented a modification of *Braess Paradox graph*, such that no strategy controlling  $\alpha \cdot r$  flow can induce  $\leq 1/\alpha$  times the optimum cost. However, we show that our main result also applies to any  $s$ - $t$  net  $G$ . We take care of the Braess graph explicitly, as a convincing example.

Joint work with A. Kaporis.