PIMS Workshop

Algorithmic Theory of Networks

March 27–29, 2015

Friday March 27

12:00 to 13:00 Lunch

- 13:00 to 14:00 Pavol Hell, Simon Fraser University
- 14:00 to 15:00 TBA

15:00 to 15:30 Coffee Break

- 15:30 to 16:30 Christian Scheideler, University of Paderborn
- 16:30 to 17:30 Jiangchuan Liu, Simon Fraser University

Saturday March 28

- 9:30 to 10:30 Sudipto Guha, University of Pennsylvania

10:30 to 11:00 Coffee Break

- 11:00 t0 13:00 Shiri Chechik, Tel-Aviv University
- 12:00 to 13:00 Tom Friedetzky, University of Durham

13:00 to 14:00 Lunch

- 14:00 to 15:00 David Kirkpatrick, University of British Columbia
- 15:00 to 16:00 Andrei Bulatov, Simon Fraser University

16:00 to 16:30 Coffee Break

- 16:30 to 17:30 Mohsen Mollanoori, University of Calgary

19:00: Dinner with all speakers

Sunday March 29

- 9:30 to 10:30 Michael Elkin, Ben-Gurion University of the Negrev

10:30 to 11:00 Coffee Break

- 11:00 t0 12:00 David Woodruff, IBM
- 12:00 to 13:00 Hoda Akbari, Simon Fraser University

13:00 to 14:00 Lunch

- 14:00 to 15:00 Peter Kling, Simon Fraser University

Hoda Akbari: Adapting Continuous Load Balancing Processes to Discrete Settings

Consider a network of processors and a set of weighted tasks arbitrarily distributed among them. Load of a node is the total weight of the tasks allocated to that node. In the neighbourhood load balancing problem, the goal is to balance the load by exchanging tasks over the edges while each node is only aware of its direct neighbours. The "continuous" load model assumes that tasks can be arbitrarily divided. Using well known analysis tools one can show that certain continuous algorithms converge to the perfectly balanced state and specify bounds on their convergence time. This is more complicated to show in the "discrete" model where tasks are atomic. Several publications propose and analyze discrete load balancing algorithms that are designed and analyzed based on continuous algorithms. In this talk, we briefly go over the work in this area and then present a discretization framework that converts a wide class of continuous load balancing processes into discrete algorithms. If the continuous algorithm converges in T rounds, then at round T of the discrete algorithm the load difference of any pair of nodes is upperbounded by $2dw_max$ where d is the degree of the graph and w_max is the maximum task weight. For general graphs, these bounds are asymptotically lower compared to the previous results. The above discretization framework is deterministic and works in a very general setting, where the tasks can have arbitrary weights and the nodes can have different speeds. Furthermore, it can be applied to a wide class of continuous processes. For the case of identical tasks, we present a randomized version of our algorithm that balances the load up to a discrepancy of $O(\sqrt{dlogn})$ provided that the initial load on every node is large enough.

The material is mainly based on the speaker's joint work with Petra Berenbrink and Thomas Sauerwald "A simple approach for adapting continuous load balancing processes to discrete settings.

Andrei Bulatov: TBA

Shiri Chechik: Approximate Nearest Neighbor Search in Metrics of Planar Graphs

In this talk we will study the problem of approximate Nearest-Neighbor Search (NNS) in graphical metrics: The task is to preprocess an edge-weighted planar graph $G = (V, E)$ on m vertices and a small "dataset" $D \subset V$ of size $n << m$, so that given a query point $q \in V$, one can quickly approximate $\text{dist}(q, D)$ (the distance from q to its closest vertex in D) and find a vertex $a \in D$ within this approximated distance.

We assume the query algorithm has access to an oracle that quickly evaluates the exact distance between any pair of vertices.

For planar graphs G with maximum degree Δ, we show how to efficiently construct a compact data structure of size $O(n\Delta)$ that quickly answers $(1 + \epsilon)$-NNS queries.

Thus, as far as NNS applications are concerned, metrics derived from bounded-degree planar graphs behave as low-dimensional metrics, even though planar metrics do not necessarily have a low doubling dimension, nor can they be embedded with low distortion into ℓ_2.

We complement our algorithmic result by matching lower bounds.

Joint work with Ittai Abraham, Robert Krauthgamer and Udi Wieder.

Michael Elkin: Distributed Symmetry Breaking

Consider an unweighted undirected graph G, and suppose that each vertex hosts a processor. The processors have unique identity numbers. They communicate over edges of G. Suppose also that the communication is synchronous, i.e., it proceeds in discrete rounds. On each round each vertex is allowed to send messages to all its neighbors. How many rounds are required to color the graph G properly in $(\Delta + 1)$ colors, where Δ is the maximum degree of the graph G? How many rounds are required to compute a maximal (under containment) independent set? These are classical questions in distributed computing. Recently significant progress was achieved in the study of both randomized and deterministic algorithms for these problems. In this talk I will overview classic and new results, and discuss some of the relevant techniques. The talk will be based on a series of papers co-authored with Leonid Barenboim, Seth Pettie, Johannes Schneider and Hsin-Hao Su.

Tom Friedetzky: Tweaking randomised load balancing approaches

We will be discussing several load balancing mechanisms based on balls-into-bins protocols and random walks. Our focus will be on making standard models more applicable, e.g., by allowing to model tasks sizes and processing speeds, or by attempting to "parallelise" inherently sequential-seeming protocols. This will be more of an overview talk light on proofs (though main ideas and techniques will be discussed).

The many authors involved in the various pieces of work will be duly mentioned during the talk.

Sudipto Guha: Dual Thresholding and Capacitated Matching

Abstract: In this talk we will see how to use dual thresholding based techniques to solve linear programs approximately and efficiently. Using this technique we will devise the first fully polynomial and near linear time approximation schemes for capacitated matching problems. The talk is based on joint work with Kook Jin Ahn and has appeared in SODA 2014.

Pavol Hell: Circular Arc Graphs

Hadwiger and Debrunner have asked (in 1960) for a characterization of circular arc graphs. While several characterizations exist, none is in terms of structural obstructions, and none could be used for a certifying recognition algorithm. We obtain such a structural obstruction characterization, and use it for a certifying recognition algorithm.

(joint with Mathew Francis and Juraj Stacho)

David Kirkpatrick: Efficient algorithms with restricted workspace: shortest paths in grid graphs using budgeted recursion.

A space-efficient algorithm for computing a minimum-weight path between two specified vertices in an arbitrary bi-directed and edge-weighted grid-graph with n vertices is described The algorithm uses only $O(\sqrt{n})$ $\Theta(\log n)$-bit words in addition to its read-only input. The path can be constructed in $O(n^{3.52})$ time, in grids with arbitrary edge weights, and $O(n^{3.02})$ time, in grids with non-negative weights.

The best previously known space-efficient algorithms for computing the weighted distance between two vertices required $O(n^{7/3})$ time with $O(n^{2/3})$ space, or $O(n^{O(1/\epsilon)})$ time with $O(n^{\epsilon}\sqrt{n})$ space, where ϵ is an arbitrarily small positive constant.

Peter Kling: Powering your Solar Cell with Duality & Variational Calculus

In a relatively recent branch of processor scheduling, classical objectives (like minimizing average response time or meeting deadlines) are combined with the need to operate devices energy-efficiently. In these scenarios, the scheduler has control over the processors' speed and must balance between a high quality of service (needing high speed and, thus, energy) and a low energy consumption. While typically motivated via the immense energy-consumption of modern data centers, similar energy-efficiency concerns arise in the deployment of sensor networks. Given that sensors often have to operate autonomously in hazardous environments, it is important to make efficient use of the available, limited energy. To prolong these system's lifetime, they're often equipped with technologies to harvest energy (e.g., solar cells).

In this talk I present recently introduced scheduling models for operating such energy harvesting systems efficiently. Using duality theory, we derive an intuitive combinatorial algorithm to optimally operate a sensor equipped with a solar cell and a battery. In another model, we use variational calculus to optimize the left over energy when operating a solar cell without a battery (selling any excess energy). The latter model introduces a technique that yields an efficient combinatorial algorithm even if the considered system is described by continuous (infinitely many) time constraints.

Jiangchuan Liu: TBA

Mohsen Mollanoori: Exploiting interference to reduce communication time in wireless networks

In wireless networks, if multiple packets are transmitted at the same time on the same channel, then the received signal is a combination of the transmitted signals. Traditionally, this is considered as a collision that requires retransmission of the packets. Modern decoding techniques, however, are able to extract useful information from multiple overlapping packets. Successive interference cancellation (SIC) and Physical-layer network coding (PNC) are two such techniques. In this talk, we examine the impact of these two techniques on the number of time slots required to achieve an objective, called Synchronization, in a wireless network consisting of several terminals and a relay node. We analyze the time complexity of a few different problems arising from this setting. In general, because of algebraic nature of PNC, problems arising from PNC can be solved in polynomial time, while, because of the combinatorial nature of SIC, problems related to SIC, are mostly NP-hard. Approximation or randomized algorithms are suggested for some NP-hard cases. Some interesting problems remain open when PNC and SIC are combined.

Christian Scheideler: Self-stabilizing overlay networks

Abstract: In this talk I will consider the problem of repairing networks from faults. More specifically, I will focus on distributed algorithms that are self-stabilizing in a sense that they can repair a network from any given set of faults as long as the network is still weakly connected. Challenges I will address in my talk are how to keep the amount of work for the repair as low as possible and how to preserve a desired functionality as the network is repaired.

David Woodruff: A Survey of Results in the Message Passing Communication Model

I'll discuss various tools for proving lower bounds in the message passing communication model, where there are k players each with an input who speak to each other in a point-to-point communication network, with the goal being to solve a function of their inputs while minimizing total communication. I'll discuss graph problems, linear-algebraic problems, and statistical problems.

This is based on joint works with Yi Li, Xiaoming Sun, Chengu Wang, and Qin Zhang