
Network Algorithms
Algorithmes des réseaux
(Org: **Evangelos Kranakis** (Carleton))

ANTHONY BONATO, Wilfrid Laurier University, Waterloo

Models of the Web Graph

One of the most widely studied real-world networks is the web graph, whose nodes represent web pages, and whose edges represent the links between pages. Another well-studied real-world network consists of protein-protein interactions in a living cell. Both networks are *self-organizing*, as each node acts as an independent agent that bases its decision on how to link to the existing network on local knowledge. Many models for self-organizing networks are *on-line*: new nodes are born over time. Hence, it is natural to consider the infinite graphs that result in the limit as time tends to infinity.

We present new results on limit graphs generated by on-line random graph models. We will characterize properties of a generalized copying model via adjacency properties, and describe self-similarity properties of limit graphs generated by on-line processes. If time permits, then we will discuss a new geometric model for the web graph.

This is joint work with Peter Cameron, Dejan Delić, and Jeannette Janssen.

MARK COATES, Department of Electrical and Computer Engineering, McGill University

Compressed Network Monitoring

This talk will describe a procedure for estimating a full set of network path metrics from a limited number of measurements. The approach exploits the strong spatial and temporal correlation observed in path-level metric data, which arises due to shared links and stationary components of the observed phenomena. We derive diffusion wavelets from the routing matrix to generate a basis in which the signals are compressible (coefficients exhibit approximately power law decay). This allows us to exploit powerful non-linear estimation algorithms that strive for sparse solutions. We demonstrate our results using measurements of end-to-end delay in the Abilene network and simulations of bit error rate in an all-photonic network. Our results show that the number of routes we need to monitor is surprisingly low (e.g., we can recover network mean end-to-end delay with 5% accuracy while monitoring only 7% of the routes).

STEFAN DOBREV, School of Information Technology and Engineering, University of Ottawa

Local Computation in Location-Aware Ad-Hoc Networks

In the same way as the spread of computer networks has motivated the study of distributed algorithms, the increase in size of these networks and appearance of wireless networks has caused significant research interest in local algorithms, in which the decision at each node is computed based only on the information contained in the node's local neighbourhood (i.e., up to k hops away, for some constant k). Local algorithms are of particular interest in wireless ad-hoc networks (e.g., sensor networks), where the dynamic and transient nature of the network requires fast adaptability and severely limits applicability of centralized and global approaches.

However, there are strong limitations on what can be computed locally, typically caused by the underlying global nature of the problem (e.g., constructing a spanning tree) or by the possibility of local symmetry. One way to deal with the problem is a relaxation of the requirements on the solution (e.g., instead of finding a spanning tree, find a sufficiently sparse subgraph), another is to consider additional structural information available in the network to break the symmetry.

In this talk, we examine how geographic information (i.e., the nodes know their GPS coordinates—are location-aware) can be exploited in ad-hoc networks modeled by unit-disc graphs. We consider the problems of locally constructing suitable (planar,

connected, with good spanner properties, of low cost) subgraphs of the underlying unit-disc graph, as well as various graph covering and colouring problems.

GEORGE KARAKOSTAS, Department of Computing & Software, McMaster University
Selfish Routing with More than One Kinds of Users

We consider the problem of characterizing user equilibria and optimal solutions for selfish routing in a given network. We extend the known models by considering users oblivious to congestion or users that can be centrally coordinated.

While in the typical selfish routing setting the users follow a strategy that minimizes their individual cost by taking into account the (dynamic) congestion due to the current routing pattern, an oblivious user ignores congestion altogether. Instead, he decides his routing on the basis of cheapest routes on a network without any flow whatsoever. These cheapest routes can be, for example, the shortest paths in the network without any flow. This model tries to capture the fact that routing tables for at least a fraction of the flow in large scale networks such as the Internet may be based on the physical distances or hops between routers alone. The phenomenon is similar to the case of traffic networks where a certain percentage of travelers base their route simply on the distances they observe on a map, without thinking (or knowing, or caring) about the delays experienced on this route due to their fellow travelers. In this work we study the price of anarchy of such networks, i.e., the ratio of the total latency experienced by the users in this setting over the optimal total latency if all users were centrally coordinated.

We also study generalization of selfish routing where the users are divided into an α fraction that are routed according to a central coordinator's routing strategy (Stackelberg strategy), and the remaining $1 - \alpha$ that determine their strategy selfishly, given the routing of the coordinated users. We analyze two such strategies from the literature, LLF and SCALE. For general topology networks, multicommodity users, and linear latency functions, we show for the first time a price of anarchy bound for SCALE, which decreases from $4/3$ to 1 as α increases from 0 to 1, and depends only on α .

MICHEL PAQUETTE, School of Computer Science, Carleton University
Communication in Networks with Positively Correlated Faults

As communication networks grow in size and complexity, they become increasingly vulnerable to component failures. To this effect, the fundamental questions of communication network reliability have been studied in the past under the assumption that components fail randomly and independently. However, it is widely accepted that faults appear with positive correlation in communication networks. In fact, empirical work has shown that this assumption is a more reasonable one than that of fault independence. To the best of our knowledge, no analytic work has been done, providing results on the feasibility and speed of communication in networks, under the assumption that components fail randomly, with positive correlation.

In this talk, we consider the problem of feasibility and time of communication in networks with dependent positively correlated faults. We propose what is, to the best of our knowledge, the first set of analytic results concerning this important aspect of network communication.

This is joint work with Evangelos Kranakis and Andrzej Pelc.

YIQIANG Q. ZHAO, School of Mathematics and Statistics, Carleton University
Tail Asymptotics and Block-Form Compensation Algorithm for Parallel Systems

Using the example of join-the-shortest-response-time model, we try to present the general idea of the block-form compensation algorithm for parallel systems and to address the connection between stationary tail asymptotics and the initial values for the implementation of the algorithm.