Ninth Summer Workshop

of the

Landelijk Netwerk Mathematische Besliskunde

NETWORKS AND OPERATIONS RESEARCH

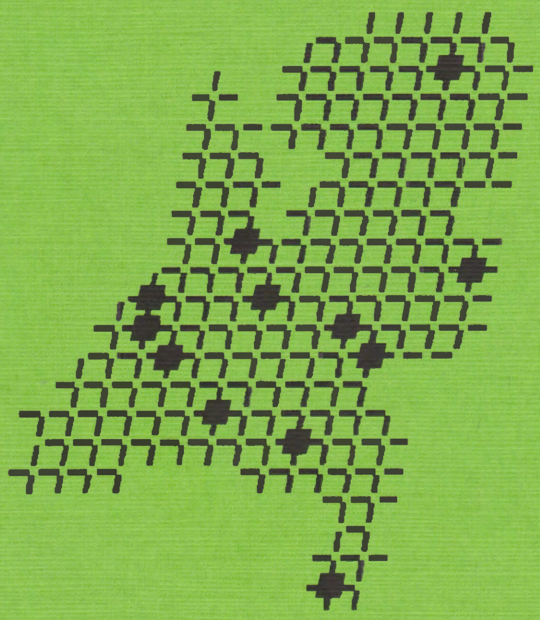

CWI, AMSTERDAM

JUNE 3-5, 1998

CWI

Ninth Summer Workshop of the Landelijk Netwerk Mathematische Besliskunde

NETWORKS AND OPERATIONS RESEARCH

June 3-5, 1998, CWI, Amsterdam

The meeting consists of ten lectures by invited speakers and two "open problem workshops".

Invited speakers

DAN BIENSTOCK (Columbia University, New York, U.S.A.)
ERIK A. VAN DOORN (University of Twente, Enschede, The Netherlands)
ANDRÁS FRANK (Eőtvős Loránd University, Budapest, Hungary)
MICHEL MANDJES (KPN Research, Leidschendam, The Netherlands)
JAMES W. ROBERTS (France Telecom-CNET, Issy-les-Moulineaux, France)

Open problem workshops

The meeting contains two "open problem workshops": afternoons during which participants will be challenged by open problems on networks; challenged for (partial) solutions if possible, but in any case to think about them, discuss them and learn!

The topics of these two sessions are: "Combinatorics of networks" and "Queueing networks".

Workshop dinner

The Workshop dinner takes place on Thursday June 4, 18.30 at the CWI.

Organization

The workshop is organized by the CWI Amsterdam and the Landelijk Netwerk Mathematische Besliskunde. The local organizers are: Onno J. Boxma, Bert Gerards, and Lex Schrijver.

Contents

Program

Wednesday June 3, 1998

10.15	*Registration*
10.45	Opening
11.00–11.45	James Roberts: Towards a multiservice network
12.00–12.45	James Roberts: Traffic engineering for stream traffic
12.45	*Lunch*
14.00–17.30	Open problem workshop: Combinatorics of networks (with an introduction by Lex Schrijver)

Thursday June 4, 1998

10.00–10.45	András Frank: Clever searching in networks
11.00–11.45	András Frank: Two-phase greedy algorithms and network connectivity
12.00–12.45	Dan Bienstock: Combinatorial network design problems: the most difficult integer programs?
12.45	*Lunch*
14.00–17.30	Open problem workshop: Queueing networks (with an introduction by Onno Boxma)
18.30	*Workshop dinner*

Friday June 5, 1998

10.00–10.45	Dan Bienstock: Modern techniques for solving large, difficult linear programs
11.00–11.45	James Roberts: Traffic engineering for elastic traffic
12.00–12.45	András Frank: How to improve the connectivity of a network
12.45	*Lunch*
14.00–14.45	Erik van Doorn: Markov-modulated fluid queues
15.00–15.45	Michel Mandjes: Optimal trajectory to overflow in fluid queues with many sources
15.45	Closing

Abstracts

Dan Bienstock
Department of Industrial Engineering and Operations Research
Columbia University
500 W 120th St.
New York, NY 10027
U.S.A.
danc@ieor.columbia.edu

COMBINATORIAL NETWORK DESIGN PROBLEMS: THE MOST DIFFICULT INTEGER PROGRAMS?

THURSDAY 12.00

Network design problems arise primarily as practical problems in telecommunications. Broadly speaking, these problems concern the assignment of capacities to components of a network (nodes and links) so that multicommodity demands can be routed between pairs of nodes. Typically, the capacity assignment part of the problem models the acquisition of expensive machinery, e.g. switches, fiber optic cables, etc. Major telecommunications companies invest large amounts of money every year engaging in this type of planning.

Some of the most difficult combinatorial problems known to researchers, such as set covering, set packing and set coloring, arise as subproblems in network design. At the same time, the routing component of the design problem is frequently complex enough to require the use of linear programming. Alas, the resulting linear programs are unexpectedly challenging.

This combination of hard combinatorial problems described by difficult linear programs yields some extremely intractable mixed-integer programs. Treated as general problems, even small instances can completely defeat the world's best integer programming solvers. Heuristics are notoriously unreliable.

However, using sophisticated algorithms and careful computation, one can obtain effective solutions to real-world problems. In this talk, we will quickly survey the history of network design problems, we will describe some of the difficulties arising in tackling them, and we will discuss some promising new techniques for their solution. One clear trend for the future is the following: if one is interested in dealing with large-scale, difficult problems, one should only implement techniques that have a solid mathematical guarantee.

MODERN TECHNIQUES FOR SOLVING LARGE, DIFFICULT LINEAR PROGRAMS

FRIDAY 10.00

Large-scale linear programs arising in the solution of several types of real-world problems, for example in transportation, and in particular, in network design, have several important features:

1. They are quite difficult (typically, they are highly primal and dual degenerate, and they have many dense rows and columns),
2. They have many more columns than rows,
3. The constraints are made up of several independent blocks (describing, for example, the routing of commodities) linked by a number of dense constrains,
4. At optimality, only a small subset of the variables is nonzero.

In recent years, theoretical advances by Plotkin, Shmoys, Tardos, and Grigoriadis and Khachiyan have resulted in polynomial-time approximation algorithms for linear programs that (roughly) have property 3 described above. Given any degree of accuracy, these algorithms will provide a solution of that degree of optimality and feasibility, after solving a polynomial number of "easy" linear programs.

We are interested in using these techniques as "crash" heuristics for a real L.P. solver, in particular in view of feature 4 above. In this talk we will present a survey of the theoretical results in this area, as well as experimental results with an implementation of the approximation algorithm for general L.P.s of type 3, but applied to combinatorial problems arising in network design.

Erik van Doorn
Faculty of Applied Mathematics
University of Twente
P.O. Box 217
7500 AE Enschede
The Netherlands
e.a.vandoorn@math.utwente.nl

MARKOV-MODULATED FLUID QUEUES FRIDAY 14.00

The model we consider is that of a fluid reservoir with infinite capacity which receives and releases fluid flows at variable rates such that the net input rate of fluid (which may be negative or positive) is uniquely determined by the state of a Markov chain evolving in the background, with the evident restriction that the content of the reservoir cannot decrease when the reservoir is empty. Some examples of (telecommunication) models fitting in this setting will be given. It will subsequently be shown how the equilibrium distribution of the content of the reservoir can be obtained when the state space of the Markov chain is finite. Then an approach to solving the problem when the state space of the Markov chain is denumerably infinite but the chain itself is a birth-death process will be described. The talk is concluded with a brief sketch of some further generalizations of the model which have recently yielded to analysis.

András Frank
Eötvös Loránd University
Department of Operations Research
Muzeum Krt. 6-8
Budapest
Hungary
frank@cs.elte.hu

CLEVER SEARCHING IN NETWORKS THURSDAY 10.00

First we briefly review two well-known techniques, depth-first and breadth-first searches, to explore the nodes of a graph, and exhibit how to use them to compute a topological ordering of acyclic digraphs, a peripherical circuit of a graph, strongly connected orientation of a graph, strongly connected components of a digraph. Second, scan-first searches and max-back (=legal=max adjacency) orderings are introduced and analyzed. We study the min-cut algorithm of Nagamochi and Ibaraki, show how to find a sparse k-connected subgraph of a k-connected graph, how to compute a simplicial node of a chordal graph.

TWO-PHASE GREEDY ALGORITHMS AND NETWORK CONNECTIVITY THURSDAY 11.00

Fulkerson described a two-phase greedy algorithm to find a minimum cost spanning arborescence and to solve the dual linear program. This was later extended for "kernel systems", a framework including the rooted edge-connectivity augmentation problem, as well. We analyze a slight extension of this model that allows us to find a minimum cost of new edges whose addition to a digraph increases its rooted node-connectivity by one. Related two-phase greedy algorithms of Faigle and Kern and of Queyranne et al. will also be outlined.

HOW TO IMPROVE THE CONNECTIVITY OF A NETWORK FRIDAY 12.00

We describe a combinatorial polynomial-time algorithm to make a $(k-1)$-connected digraph k-connected by adding a minimum number of new edges, or more generally, a minimum cost of new edges with respect to a node-cost induced cost function. A related algorithm will also be exhibited for a generalization of a theorem of E. Győri and for its weighted extension by A. Lubiw to find a minimum set of (weighted) generators of a family of subpaths of a circuit.

Michel Mandjes
KPN Research
Leidschendam
The Netherlands
m.r.h.mandjes@research.kpn.com

OPTIMAL TRAJECTORY TO OVERFLOW IN FLUID QUEUES WITH MANY SOURCES
FRIDAY 15.00

Joint work with Ad Ridder (Free University Amsterdam)

We analyse the deviant behavior of a queue fed by a large number of traffic streams, which models an ATM switch. In particular, we explicitly give the most likely trajectory (or 'optimal path') to buffer overflow, by applying large deviations techniques. This is done for a broad class of sources, consisting of Markov fluid sources and periodic sources. Apart from a number of ramifications of this result, we present guidelines for the numerical evaluation of the optimal path.

We consider a finite-buffer queue that is fed by a large number of traffic sources. The sources are either all of Markov Fluid type or all deterministic (periodic) on/off type. These sources generate fluid streams into the buffer which is emptied at a constant rate. Such queueing models attract nowadays much attention in performance evaluation of ATM networks.

Our study concerns the rare event of a buffer overflow, more specifically, the most likely trajectory towards overflow. A trajectory or path f gives at time t the fraction $f_i(t)$ of sources that is in state i. Within the class of all feasible paths there is a set of paths F leading to buffer overflow. In this set there is a path f^* which the sources follow with 'overwhelming probability' for reaching overflow. Formally, f^* minimizes a cost functional J on the set F. The cost functional has been identified by Weiss [1], who also solves the optimization for the optimal trajectory in case all sources are homogeneous exponential on/off.

The contribution of our paper is that we find explicitly the optimal trajectory for general heterogeneous Markov fluid sources. The approach we follow is to apply large deviations techniques—similar to Botvich and Duffield [2] or Simonian and Guibert [3]—and to prove optimality as in [1]. In case of deterministic sources, we introduce the notion of optimal trajectory towards overflow and find explicitly its expression. Finally, these results are illustrated numerically.

References

[1] A. Weiss. A new technique of analyzing large traffic systems. *Advances of Applied Probability*, 18: 506 – 532, 1986.
[2] D.D. Botvich and N.G. Duffield. Large deviations, the shape of the loss curve, and economies of scale in large multiplexers. *Queueing Systems*, 20: 293 – 320, 1995.
[3] A. Simonian and J. Guibert. Large deviations approximation for fluid queues fed by a large number of on/off sources. *IEEE Journal of Selected Areas in Communications*, 13: 1017 – 1027, 1995.

James Roberts

France Telecom-CNET
38 Rue du General Leclerc
Cedex 9, Issy-les-Moulineaux 92794
France
`james.roberts@cnet.francetelecom.fr`

TOWARDS A MULTISERVICE NETWORK — WEDNESDAY 11.00

In this talk we discuss proposed solutions for a multiservice network coming respectively from ATM and IP communities. We start with a discussion of user requirements identifying two broad traffic categories: "stream" and "elastic". Stream flows are generated by voice and video applications and generally require the network to preserve time integrity. Elastic flows, on the other hand, have no intrinsic rate or duration; their quality of service is ultimately measured by realized throughput. We discuss the adequacy of proposed service classes with respect to the quality of service guarantees required by these two types of traffic. In particular, we highlight the impact of charging on the choice of service model and, indeed, on the feasibility of quality of service guarantees.

TRAFFIC ENGINEERING FOR STREAM TRAFFIC — WEDNESDAY 12.00

Stream traffic flows (for voice and video) are suitable for open loop or preventive control based on statistical multiplexing. We consider here the performance of a statistical multiplexer, outlining evaluation methods and indicating qualitative results. The predictability of multiplexing performance determines the feasibility of traffic control options, notably for connection admission control. We argue in favour of a control option for stream traffic based on "bufferless" multiplexing with CAC based on measured traffic rather than a priori traffic descriptors. Network dimensioning for stream traffic can then be performed to ensure limited blocking probabilities using multirate extensions of the Erlang model.

TRAFFIC ENGINEERING FOR ELASTIC TRAFFIC — FRIDAY 11.00

There is currently little established traffic engineering theory for networks performing reactive control to share bandwidth between elastic traffic flows (like the Internet). Assuming a fixed set of flows, we first consider desirable bandwidth sharing objectives (max-min or proportional fairness) and the algorithms by which they can be achieved. We then examine performance under random traffic assumptions, observing notably that the usual notions of fairness or weighted fairness do not lead to optimal performance from the user or the network. We argue for the introduction of admission control for elastic flows to ensure minimal throughput guarantees with surplus bandwidth shared to minimize blocking probabilities. We conclude with the outline of a simple integrated services service model fulfilling the requirements identified in the previous three talks.

LNMB a'98
works.tex
worksh.tex

LNMB Workshop, June 4, 1998: Open queueing problems

Problem 1
Consider a polling system, consisting of two queues Q_1, Q_2 and a single server. Customers arrive at the two queues according to independent Poisson processes, with rates $\lambda_1 = \lambda_2 = 1$. The service times at each queue are independent and exponentially distributed, with mean 1/3. The server serves one customer at Q_1 (when there is somebody to be served), then switches to Q_2 (this takes a constant time of length 0.1). At Q_2 he serves exactly those customers who were already present at the moment of his own arrival (this rule is called 'gated'). Then he switches back to Q_1 (this takes a constant time of length 0.1). Note that according to the above rules, the server keeps switching between the two queues in an empty system.
Qu.: what is the mean waiting time of a customer in Q_1? And what is the Laplace-Stieltjes transform of the waiting time distribution? These are open problems, that are of considerable importance in various communication settings. Good approximations are already valuable.

Problem 2
Consider two single-server queues Q_1 and Q_2 in series. Each server uses the FCFS discipline. Q_1 is an M/D/1 queue with service times equal to 1 and arrival intensity 0.5. The service times in Q_2 are exponentially distributed with mean 0.5.
Qu.: what is the mean total sojourn time of a customer? And what is the Laplace-Stieltjes transform of the sojourn time distribution?

Problem 3
Consider a closed system of two single-server queues Q_1 and Q_2. The system contains N customers. Each server uses the FCFS discipline. Q_1 is a $./D/1$ queue with service times equal to 1. Q_2 is a $./M/1$ queue with service times that are exponentially distributed of mean a.
Qu.: what is the mean total sojourn time of a customer in (first) Q_1 and (then) Q_2? And what is the Laplace-Stieltjes transform of this sojourn time distribution?

LNMB Workshop, June 4, 1998: Open queueing problems

Problem 4

In the CWI-canteen, one can distinguish two types of customers: those who want a hot meal (type-1), and those who don't (type-2). Both types arrive according to independent Poisson processes, with rate $\lambda_1 = 1$ per minute and $\lambda_2 = 3$ per minute. There is one server for the hot meals. The service times here are independent, exponentially distributed with mean 0.8 minute. There are two checkout counters; service times there are exponentially distributed with mean 0.2 minute.

Some time ago, the canteen lay-out has been changed. Let us investigate the effect of this change.

Figure 1 represents the old canteen lay-out. 'Zelfbed.' represents self-service (taking milk, cheese, etc.; this requires 0.5 minute (let's say a non-random period of time)). Customers used to choose checkout counter 1 with probability $q = 5/8$ and counter 2 with probability $1 - q = 3/8$.

Qu.: determine the mean total sojourn time of an arbitrary customer before he/she has passed the checkout counter.

Figure 2 represents the new canteen lay-out. There are separate lanes for type-1 and type-2 customers. Self-service still requires 0.5 minute. A fraction α of the type-2 customers moves to counter 1.

Qu.: Determine the mean total sojourn time of an arbitrary customer. Which α minimizes the mean total sojourn time of an arbitrary customer?

Problem 5

Develop a queueing model (customers, arrival processes, service times, queues, servers, service disciplines) for the traffic light system at the crossing of Middenweg and Kruislaan. Furthermore, suggest a way to analyse and optimise this system. If you want to restrict yourself to cars, that is OK.

Note: in each of the four roads, there are two lanes: in the right lane, a car may turn right or drive straight ahead, and in the left lane it may turn left or drive straight ahead.

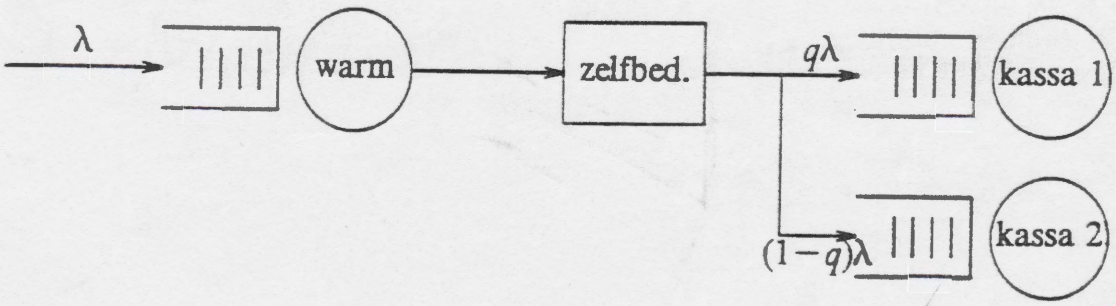

Figuur 1

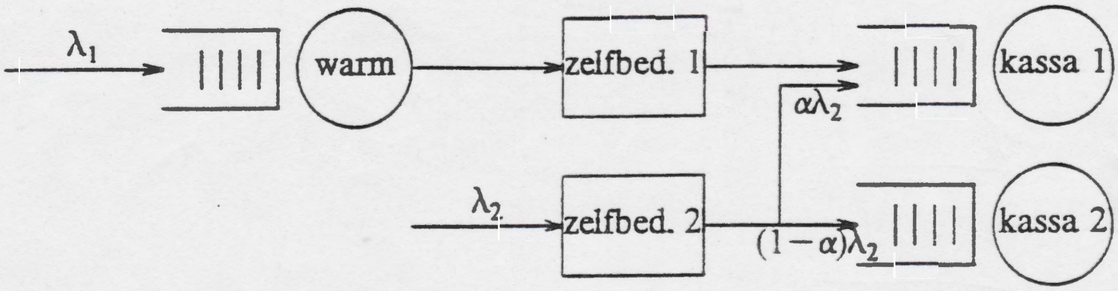

Figuur 2