

GERAD

Groupe d'études et de recherche en analyse des décisions

PERFORMANCE EVALUATION AND PLANNING METHODS FOR THE NEXT GENERATION INTERNET

Edited by

André Girard

Brunilde Sansò

Felisa Vázquez-Abad

 Springer

PERFORMANCE EVALUATION AND PLANNING METHODS FOR THE NEXT GENERATION INTERNET

GERAD 25th Anniversary Series

- **Essays and Surveys in Global Optimization**
Charles Audet, Pierre Hansen, and Gilles Savard, editors
- **Graph Theory and Combinatorial Optimization**
David Avis, Alain Hertz, and Odile Marcotte, editors
- **Numerical Methods in Finance**
Hatem Ben-Ameur and Michèle Breton, editors
- **Analysis, Control and Optimization of Complex Dynamic Systems**
El-Kébir Boukas and Roland Malhamé, editors
- **Column Generation**
Guy Desaulniers, Jacques Desrosiers, and Marius M. Solomon, editors
- **Statistical Modeling and Analysis for Complex Data Problems**
Pierre Duchesne and Bruno Rémillard, editors
- **Performance Evaluation and Planning Methods for the Next Generation Internet**
André Girard, Brunilde Sansò, and Félisa Vázquez-Abad, editors
- **Dynamic Games: Theory and Applications**
Alain Haurie and Georges Zaccour, editors
- **Logistics Systems: Design and Optimization**
André Langevin and Diane Riopel, editors
- **Energy and Environment**
Richard Loulou, Jean-Philippe Waaub, and Georges Zaccour, editors

PERFORMANCE EVALUATION AND PLANNING METHODS FOR THE NEXT GENERATION INTERNET

Edited by

ANDRÉ GIRARD

GERAD and INRS-Énergie, Matériaux et Télécommunications

BRUNILDE SANSÒ

GERAD and École Polytechnique de Montréal

FELISA VÁZQUEZ-ABAD

GERAD and Université de Montréal



Springer

André Girard
GERAD & INRS-Telecommunications
Montréal, Canada

Brunilde Sansò
GERAD & École Polytechnique de Montréal
Montréal, Canada

Felisa Vázquez-Abad
GERAD and Université de Montréal
Montréal, Canada

Library of Congress Cataloging-in-Publication Data

A C.I.P. Catalogue record for this book is available
from the Library of Congress.

ISBN-10: 0-387-25550-8 ISBN 0-387-25551-6 (e-book) Printed on acid-free paper.
ISBN-13: 978-0387-25550-7

© 2005 by Springer Science+Business Media, Inc.

All rights reserved. This work may not be translated or copied in whole or in part without the written permission of the publisher (Springer Science + Business Media, Inc., 233 Spring Street, New York, NY 10013, USA), except for brief excerpts in connection with reviews or scholarly analysis. Use in connection with any form of information storage and retrieval, electronic adaptation, computer software, or by similar or dissimilar methodology now known or hereafter developed is forbidden.

The use in this publication of trade names, trademarks, service marks and similar terms, even if they are not identified as such, is not to be taken as an expression of opinion as to whether or not they are subject to proprietary rights.

Printed in the United States of America.

9 8 7 6 5 4 3 2 1

SPIN 11053163

springeronline.com

Foreword

GERAD celebrates this year its 25th anniversary. The Center was created in 1980 by a small group of professors and researchers of HEC Montréal, McGill University and of the École Polytechnique de Montréal. GERAD's activities achieved sufficient scope to justify its conversion in June 1988 into a Joint Research Centre of HEC Montréal, the École Polytechnique de Montréal and McGill University. In 1996, the Université du Québec à Montréal joined these three institutions. GERAD has fifty members (professors), more than twenty research associates and post doctoral students and more than two hundreds master and Ph.D. students.

GERAD is a multi-university center and a vital forum for the development of operations research. Its mission is defined around the following four complementarily objectives:

- The original and expert contribution to all research fields in GERAD's area of expertise;
- The dissemination of research results in the best scientific outlets as well as in the society in general;
- The training of graduate students and post doctoral researchers;
- The contribution to the economic community by solving important problems and providing transferable tools.

GERAD's research thrusts and fields of expertise are as follows:

- Development of mathematical analysis tools and techniques to solve the complex problems that arise in management sciences and engineering;
- Development of algorithms to resolve such problems efficiently;
- Application of these techniques and tools to problems posed in related disciplines, such as statistics, financial engineering, game theory and artificial intelligence;
- Application of advanced tools to optimization and planning of large technical and economic systems, such as energy systems, transportation/communication networks, and production systems;
- Integration of scientific findings into software, expert systems and decision-support systems that can be used by industry.

One of the marking events of the celebrations of the 25th anniversary of GERAD is the publication of ten volumes covering most of the Center's research areas of expertise. The list follows: **Essays and Surveys in Global Optimization**, edited by C. Audet, P. Hansen and G. Savard; **Graph Theory and Combinatorial Optimization**,

edited by D. Avis, A. Hertz and O. Marcotte; **Numerical Methods in Finance**, edited by H. Ben-Ameur and M. Breton; **Analysis, Control and Optimization of Complex Dynamic Systems**, edited by E.K. Boukas and R. Malhamé; **Column Generation**, edited by G. Desaulniers, J. Desrosiers and M.M. Solomon; **Statistical Modeling and Analysis for Complex Data Problems**, edited by P. Duchesne and B. Rémillard; **Performance Evaluation and Planning Methods for the Next Generation Internet**, edited by A. Girard, B. Sansò and F. Vázquez-Abad; **Dynamic Games: Theory and Applications**, edited by A. Haurie and G. Zaccour; **Logistics Systems: Design and Optimization**, edited by A. Langevin and D. Riopel; **Energy and Environment**, edited by R. Loulou, J.-P. Waaub and G. Zaccour.

I would like to express my gratitude to the Editors of the ten volumes, to the authors who accepted with great enthusiasm to submit their work and to the reviewers for their benevolent work and timely response. I would also like to thank Mrs. Nicole Paradis, Francine Benoît and Louise Letendre and Mr. André Montpetit for their excellent editing work.

The GERAD group has earned its reputation as a worldwide leader in its field. This is certainly due to the enthusiasm and motivation of GERAD's researchers and students, but also to the funding and the infrastructures available. I would like to seize the opportunity to thank the organizations that, from the beginning, believed in the potential and the value of GERAD and have supported it over the years. These are HEC Montréal, École Polytechnique de Montréal, McGill University, Université du Québec à Montréal and, of course, the Natural Sciences and Engineering Research Council of Canada (NSERC) and the Fonds québécois de la recherche sur la nature et les technologies (FQRNT).

Georges Zaccour
Director of GERAD

Avant-propos

Le Groupe d'études et de recherche en analyse des décisions (GERAD) fête cette année son vingt-cinquième anniversaire. Fondé en 1980 par une poignée de professeurs et chercheurs de HEC Montréal engagés dans des recherches en équipe avec des collègues de l'Université McGill et de l'École Polytechnique de Montréal, le Centre comporte maintenant une cinquantaine de membres, plus d'une vingtaine de professionnels de recherche et stagiaires post-doctoraux et plus de 200 étudiants des cycles supérieurs. Les activités du GERAD ont pris suffisamment d'ampleur pour justifier en juin 1988 sa transformation en un Centre de recherche conjoint de HEC Montréal, de l'École Polytechnique de Montréal et de l'Université McGill. En 1996, l'Université du Québec à Montréal s'est jointe à ces institutions pour parrainer le GERAD.

Le GERAD est un regroupement de chercheurs autour de la discipline de la recherche opérationnelle. Sa mission s'articule autour des objectifs complémentaires suivants :

- la contribution originale et experte dans tous les axes de recherche de ses champs de compétence ;
- la diffusion des résultats dans les plus grandes revues du domaine ainsi qu'auprès des différents publics qui forment l'environnement du Centre ;
- la formation d'étudiants des cycles supérieurs et de stagiaires post-doctoraux ;
- la contribution à la communauté économique à travers la résolution de problèmes et le développement de coffres d'outils transférables.

Les principaux axes de recherche du GERAD, en allant du plus théorique au plus appliqué, sont les suivants :

- le développement d'outils et de techniques d'analyse mathématiques de la recherche opérationnelle pour la résolution de problèmes complexes qui se posent dans les sciences de la gestion et du génie ;
- la confection d'algorithmes permettant la résolution efficace de ces problèmes ;
- l'application de ces outils à des problèmes posés dans des disciplines connexes à la recherche opérationnelle telles que la statistique, l'ingénierie financière, la théorie des jeux et l'intelligence artificielle ;
- l'application de ces outils à l'optimisation et à la planification de grands systèmes technico-économiques comme les systèmes énergétiques, les réseaux de télécommunication et de transport, la logistique et la distributive dans les industries manufacturières et de service ;

- l'intégration des résultats scientifiques dans des logiciels, des systèmes experts et dans des systèmes d'aide à la décision transférables à l'industrie.

Le fait marquant des célébrations du 25^e du GERAD est la publication de dix volumes couvrant les champs d'expertise du Centre. La liste suit : **Essays and Surveys in Global Optimization**, édité par C. Audet, P. Hansen et G. Savard ; **Graph Theory and Combinatorial Optimization**, édité par D. Avis, A. Hertz et O. Marcotte ; **Numerical Methods in Finance**, édité par H. Ben-Ameur et M. Breton ; **Analysis, Control and Optimization of Complex Dynamic Systems**, édité par E.K. Boukas et R. Malhamé ; **Column Generation**, édité par G. Desaulniers, J. Desrosiers et M.M. Solomon ; **Statistical Modeling and Analysis for Complex Data Problems**, édité par P. Duchesne et B. Rémillard ; **Performance Evaluation and Planning Methods for the Next Generation Internet**, édité par A. Girard, B. Sansò et F. Vázquez-Abad ; **Dynamic Games : Theory and Applications**, édité par A. Haurie et G. Zaccour ; **Logistics Systems : Design and Optimization**, édité par A. Langevin et D. Riopel ; **Energy and Environment**, édité par R. Loulou, J.-P. Waaub et G. Zaccour.

Je voudrais remercier très sincèrement les éditeurs de ces volumes, les nombreux auteurs qui ont très volontiers répondu à l'invitation des éditeurs à soumettre leurs travaux, et les évaluateurs pour leur bénévolat et ponctualité. Je voudrais aussi remercier Mmes Nicole Paradis, Francine Benoît et Louise Letendre ainsi que M. André Montpetit pour leur travail expert d'édition.

La place de premier plan qu'occupe le GERAD sur l'échiquier mondial est certes due à la passion qui anime ses chercheurs et ses étudiants, mais aussi au financement et à l'infrastructure disponibles. Je voudrais profiter de cette occasion pour remercier les organisations qui ont cru dès le départ au potentiel et à la valeur du GERAD et nous ont soutenus durant ces années. Il s'agit de HEC Montréal, l'École Polytechnique de Montréal, l'Université McGill, l'Université du Québec à Montréal et, bien sûr, le Conseil de recherche en sciences naturelles et en génie du Canada (CRSNG) et le Fonds québécois de la recherche sur la nature et les technologies (FQRNT).

Georges Zaccour
Directeur du GERAD

Contents

Foreword	v
Avant-propos	vii
Contributing Authors	xi
Preface	xiii
1	
Design of IP Networks with End-to-End Performance Guarantees <i>I. Atov and R.J. Harris</i>	1
2	
Design of IP Virtual Private Networks under End-to-end QoS Constraints <i>E.C.G. Wille, M. Mellia, E. Leonardi, and M.A. Marsan</i>	35
3	
Design of Protected Working Capacity Envelopes Based on p -Cycles: An Alternative Framework for Survivable Automated Lightpath Provisioning <i>G. Shen and W.D. Grover</i>	63
4	
Network Traffic Engineering with Varied Levels of Protection in the Next Generation Internet <i>S. Srivastava, S.R. Thirumalasetty, and D. Medhi</i>	99
5	
Balancing Traffic Flows in Resilient Packet Rings <i>P. Kubat and J. MacGregor Smith</i>	125
6	
Game-Theoretic Resource Pricing for the Next Generation Internet <i>B.M. Ninan and M. Devetsikiotis</i>	141
7	
A New Approach to Policy-Based Routing in the Internet <i>B.R. Smith and J.J. Garcia-Luna-Aceves</i>	165
8	
Advanced Methods for the Estimation of the Origin Destination Traffic Matrix <i>S. Vaton, J.S. Bedo and A. Gravey</i>	189
9	
Energy and Cost Optimizations in Wireless Sensor Networks: A Survey <i>V. Mhatre and C. Rosenberg</i>	227

	<i>NEXT GENERATION INTERNET</i>
x	
10	
Duality-Based TCP Congestion Control with Error Analysis	249
<i>M. Mehyar, D. Spanos, and S.H. Low</i>	
11	
Fast Algorithmic Solutions to Multi-dimensional Birth-Death Processes with Applications to Telecommunication Systems	269
<i>L. D. Servi</i>	
12	
A New Paradigm for On-Line Management of Communication Networks with Multiplicative Feedback Control	297
<i>H. Yu and C.G. Cassandras</i>	
13	
Comparing Locality of Reference – Some Folk Theorems for the Miss Rate and the Output of Caches	333
<i>A.M. Makowski and S. Vanichpun</i>	

Contributing Authors

IRENA ATOV
Swinburne University of Technology,
Australia
iatov@swin.edu.au

JEAN-SÉBASTIEN BEDO
École Polytechnique, Palaiseau, France
jean-sebastien.bedo@polytech-
nique.fr

CHRISTOS G. CASSANDRAS
Boston University, USA
cgc@bu.edu

MICHAEL DEVETSIKIOTIS
North Carolina State University, USA
mdevets@ncsu.edu

JOSE JOAQUIN GARCIA-LUNA-ACEVES
University of California, Santa Cruz,
USA
jj@cse.ucsc.edu

ANNIE GRAVEY
ENST Bretagne, France
annie.gravey@enst-bretagne.fr

WAYNE D. GROVER
TRLabs and University of Alberta,
Canada
grover@trlabs.ca

RICHARD J. HARRIS
Massey University, New Zealand
r.j.harris@massey.ac.nz

PETER KUBAT
Verizon Laboratories, USA
peter.kubat@verizon.com

EMILIO LEONARDI
Politecnico di Torino, Italy
leonardi@polito.it

STEVEN H. LOW
California Institute of Technology,
USA
slow@caltech.edu

JAMES MACGREGOR SMITH
University of Massachusetts, USA
jmsmith@ecs.umass.edu

ARMAND M. MAKOWSKI
University of Maryland
armand@isr.umd.edu

MARCO AJMONE MARSAN
Politecnico di Torino, Italy
ajmone@polito.it

DEEP MEDHI
University of Missouri–Kansas City,
USA
dmedhi@umkc.edu

MORTADA MEHYAR
California Institute of Technology, USA
morr@caltech.edu

MARCO MELLIA
Politecnico di Torino, Italy
mellia@prezzemolo.polito.it

VIVEK PRAKASH MHATRE
Purdue University, USA
mhatre@ecn.purdue.edu

BOBBY M. NINAN
North Carolina State University, USA
bmninan@ncsu.edu

CATHERINE ROSENBERG
University of Waterloo, Canada
cath@ece.uwaterloo.ca

L. D. SERVI
MIT Lincoln Laboratory, USA
servi@ll.mit.edu

GANGXIANG SHEN
TRLabs and University of Alberta,
Canada
gshen@trlabs.ca

BRADLEY R. SMITH
University of California, Santa Cruz,
USA
brad@soe.ucsc.edu

DEMETRI SPANOS
California Institute of Technology,
USA
demetri@caltech.edu

SHEKHAR SRIVASTAVA
University of Missouri–Kansas City,
USA
shekh@conrel.sice.umkc.edu

SRINIVASA RAO THIRUMALASETTY
Ciena Corporation, USA
sthiruma@ciena.com

SARUT VANICHPUN
University of Maryland
sarut@eng.umd.edu

SANDRINE VATON
ENST Bretagne, France
sandrine.vaton@enst-bretagne.fr

EMILIO C.G. WILLE
Centro Federal de Educação Tecnológica
do Paraná, Brazil
wille@prezzemolo.polito.it

HAINING YU
Boston University, USA
fernyu@bu.edu

Preface

Optimization techniques have been used for a long time in the planning of telecommunication networks. There is abundant work on the design of telephone and transmission networks going at least a half century back. The recent evolution towards an integrated, multi-service network based on the Internet and the IP protocol has occurred in a very different way and without much recourse to the traditional traffic engineering methods. Because the current core network is so much over-provisioned and all the services are operating on a best-effort service model, there has been no real need for sophisticated planning modeling and analysis.

This situation is bound to change for a number of reasons. New services such as voice are being introduced with strong requirements for definite Quality of Service, the cost of over-provisioning private networks is becoming an issue and wireless access cannot be over-provisioned at all since the bandwidth is limited by the available spectrum. For all these reasons, we believe that all the classical problems of telecommunication network design will become more essential in the coming years. One such problem is network design which in turn is based on performance evaluation which is an essential element of all design algorithms.

This trend is well illustrated by the contents of this book. We find a large number of areas where significant work is being done to address the issue of network design in the context of the new IP-based networks. The topics selected here will give the reader some idea of what is going on but is far from exhaustive for obvious space and time limitations.

The design of IP networks will have to take into account the requirements of many applications for guaranteed performance. This problem is examined in “Design of IP Networks with End-to-end Performance Guarantees” by Atov and Harris. Their model takes into account the proposed QoS standards and can handle multiple QoS requirements. The numerical method based on multicommodity flows can compute networks of realistic size.

A similar problem of network design is the subject of “Design of IP Virtual Private Networks under End-To-End QoS Constraints” by Wille, Mellia, Leonardi and Ajmone Marsan. This time, the network to be designed is a Virtual Private Network but the model takes into account both the transport and the network layer performances. The design itself is then done at the IP layer while offering all the required QoS guarantees at the transport layer.

The concept of resilience or reliability is gaining importance as another aspect of the Quality of Service that will have to be provided by the Internet. The notion of survivability by using protection cycles is examined in “Design of Protected Working Capacity Envelopes Based on P-Cycles: an Alternative Framework for Survivable Automated Light-path Provisioning” by Shen and Grover when applied to the underlying optical transport infrastructure. The model integrates both the IP and the transport layers to provide a unified design methodology for the capacity allocation, service provisioning and reserve network.

A similar problem of network design where the notion of QoS is extended to include reliability is examined in “Network Traffic Engineering with Varied Levels of Protection in the Next Generation Internet” by Srivastava, Thirumalasetty and Medhi. In addition to traffic engineering, reliability is taken care of via a number of protection levels. The model is based on protection cycles and is solved via heuristics.

The notion of resilience in transmission networks, an essential complement to IP-level techniques, is studied in “Balancing Traffic Flows in Resilient Packet Rings” by Kubat and Smith. Here fast restoration techniques based on ring topologies are used in conjunction with Ethernet technology to provide a robust network. The model optimizes the allocation of traffic to the ring directions both for deterministic and stochastic demands either via an IP formulation or using heuristics for large cases.

Pricing and routing are two issues that can be closely tied to the management of QoS. In “Game-Theoretic Resource Pricing for the Next Generation Internet”, Ninan and Devetsikiotis use pricing and billing to manage the bandwidth allocated to users competing for a share of the network resources. They show applications of their model to a variety of networks.

In addition to a thorough review of existing QoS-based routing, the chapter “A New Approach to Policy-Based Routing in the Internet” by Smith and Garcia-Luna-Aceves describes a new routing algorithm based on distributed label-swapping that can support QoS more efficiently than the present techniques.

Network design is based on a good estimation of the traffic matrix. This problem is examined in the context of IP networks in the paper “Advanced Methods for the Estimation of the Origin Destination Traffic Matrix” by Vatou, Bedo and Gravey. Current tools such as SNMP allow traffic measurement only on individual links. Statistical methods are used to take into account the time variation of traffic measured on the network links to construct an estimate of the end-to-end demands.

Wireless networks are growing even more rapidly than the Internet and bring up a new set of design problems. The recent development of sensor networks is a case in point where the management of the energy budget of the batteries in the devices is a dominant issue. A thorough survey of the routing and design problems in this new context is the subject of “Energy and Cost Optimizations in Wireless Sensor Networks: A Survey” by Mhatre and Rosenberg.

If they are to provide the QoS required by applications, network operators must be able to control congestion whenever it occurs in the network. An approach to congestion control based on the maximization of utility is used in “Duality-Based TCP Congestion Control with Error Analysis” by Mehyar, Spanos and Low. The model provides a unifying framework for a large number of congestion control algorithms. The authors also show that even in the presence on imperfect information, the control will converge to a region close to the optimal operating point if perfect information were available.

Network design is based on decomposition techniques where the performance of output queues has to be computed separately before being recombined into a network performance measure. Thus queuing analysis is a fundamental component of network planning. With multi-service networks, this analysis requires the solution of large birth-death systems. An example of an efficient solution algorithm is given in “Fast Algorithmic Solutions to Multi-Dimensional Birth-Death Processes with Applications to Telecommunication Systems” by Servi where a new class of solution algorithms is presented that can be applied to systems of arbitrary dimensions.

In “A New Paradigm for On-Line Management of Communication Networks with Multiplicative Feedback Control” by Yu and Cassandras, Stochastic Flow models are used for control. The model obviates the need to solve difficult queuing analysis problems by using fluid models. It provides an estimation of the gradients of the performance which in turn can be used to adjust the flow rates via a multiplicative feedback algorithm.

The Internet has spawned a number of new applications with previously unknown operating characteristics. Caching techniques are often used to reduce the delay in accessing databases. Many such techniques have been designed based on empirical knowledge. The validity of this knowledge is formalized and evaluated in “Comparing Locality of Reference — Some Folk Theorems for the Miss Rate and the Output of Caches” by Makowski and Vanichpun where the authors model the op-

eration of caches and examine the efficiency of various cache replacement policies.

The wide range of topics covered in this book is a witness to the diversity of problems that have to be faced if we want to design the Next Generation IP Networks to meet their expected performance. This is a rich field where optimization techniques can provide significant gains.

ANDRÉ GIRARD
BRUNILDE SANSÒ
FELISA VÁZQUEZ-ABAD

Chapter 1

DESIGN OF IP NETWORKS WITH END-TO-END PERFORMANCE GUARANTEES

Irena Atov
Richard J. Harris

Abstract In this paper, we examine the issues that surround IP network design with quality of service (QoS) guarantees and propose a new network design methodology. The proposed network design model takes account of the new QoS technologies (i.e., DiffServ/MPLS) and allows for multiple delay constraints so that guaranteed performance can be achieved for each of the traffic classes. After discussing the most crucial planning issues that must be addressed when QoS mechanisms are used in an IP-based network, a non-linear multicommodity optimisation problem is formulated and heuristics for its approximate solutions are described. The network design model is evaluated in terms of accuracy and scalability for each of the main components that the model employs. The computational results for each of the building blocks demonstrate that realistic size problems can be solved with the proposed method.

1. Introduction

The development of various Internet technologies, such as DiffServ and MPLS, has enabled support for various traffic classes with different QoS requirements on an integrated IP network (Wang, 2001). Now with the inclusion of QoS considerations, the paradigm for network design and planning must change to include multiple delay constraints so that differentiated performance can be achieved for the various traffic classes. The delay QoS metric is additive (i.e., QoS along a path is a sum of the QoS of its constituent links) and such a QoS class may cover a delay bound or random variation of delay (jitter).

It is traditionally difficult to provide network design to meet varying QoS constraints. Assuming that the network topology is given, the ob-

jective of network design is to determine link capacities combined with traffic routing, such that the total network cost is minimised while meeting demands and QoS constraints for each of the traffic classes. With the introduction of Asynchronous Transfer Mode (ATM) broadband networks in the late 80s a lot of research was triggered in this area. However, the network design models developed for ATM use a loss-based approach. That is, network planning-design is achieved so that the blocking (loss) probabilities of various types of traffic (classes) remain below a specified threshold, see for example Liang and Ross (1999); Pua (1999). These models exploit the connection-oriented nature of ATM and transform the multi-level traffic problem into a multirate circuit-switched problem by using the notion of equivalent bandwidth (Guerin et al., 1991). However, these multirate loss models do not present suitable tools for the design of IP-based broadband networks. In IP-based broadband networks (i.e., DiffServ/MPLS), due to the concept of class aggregates and static resource reservation, network design is primarily concerned with developing performance guarantees in terms of packet delay or packet delay variation for the various service types (i.e., class aggregates). In addition, the conventional methods for capacity planning of IP networks are limited, in that they only consider best effort service, or else a single delay constraint for all traffic (Gavish and Neuman, 1989). Thus, there is a need for new methods for capacity planning and design that take account of technologies and mechanisms that enable QoS in IP networks and, thus, allow for multiple delay constraints, so that guaranteed performance can be achieved for each of the traffic classes. This paper investigates the issues surrounding IP network design with QoS and proposes a new design methodology which would be applicable to such networks. Since the proposed design model deals in a unified way with both the flow and the capacity assignment issues, we shall refer to it in the following as a capacity and flow assignment problem (CFA problem) as it is of the same generic form as early ARPANET models identified by Kleinrock (1975).

There are major challenges involved in the development of an IP network design methodology that supports guaranteed services on an end-to-end basis. The technologies that provide QoS introduce new constraints and require that certain features be addressed by any generic design methodology. In Section 2, we first discuss the most crucial planning issues that must be addressed when QoS mechanisms are used in an IP-based network. Then we outline a network model and a cost model, which define the set of underlying assumptions used in the development of the proposed IP network design model with end-to-end performance guarantees. Section 3 provides the notation and formu-

lation of the mathematics programming problem. The CFA problem is formulated as a non-linear multicommodity optimisation problem, which is hard to solve mathematically to optimum for practical sized networks. To be able to efficiently solve large problem instances, we develop a framework for the solution of the network design problem which employs a heuristic approach. Section 4 describes the disaggregation of the problem into simpler optimisation problems (components) that form the basis for a heuristic solution to the original problem. Finally, Sections 5 and 6 summarize computational results and provide concluding remarks.

2. QoS mechanisms and implications for network planning

QoS mechanisms that are being employed in IP-based core networks can be generally classified into three main categories: traffic control, resource management and traffic engineering. In the following, we briefly overview these categories and discuss the implications that have to be considered for the planning process.

Traffic control. Traffic control encompasses all mechanisms for handling and forwarding of packets within the edge and the core routers of the network. These mechanisms include: traffic classification and aggregation, scheduling and active queue management. Typically, based on the type of scheduling mechanism deployed at the routers, which can range from a simple priority queueing to a bandwidth allocation scheme (WFQ, WF2Q, WRR Floyd and Jacobson, 1995), the network provider can offer two types of service quality. The first type is prioritised service, where certain class of traffic receives priority over others as it is processed and routed over the network. The second type is guaranteed service, where the traffic classes are guaranteed certain share of resources, for example bandwidth or a given performance level e.g., delay. It is the latter meaning that we consider in this paper, as it represents more common implementation of DiffServ in the current service offerings. The use of bandwidth allocation type QoS queueing mechanism implies that *fixed bandwidth partitioning*, i.e., predictable capacity assignment to individual classes can be assumed for the network design process. Since the DiffServ technology is based on the aggregation of individual flows into classes at the ingress of the network and on provisioning of QoS to the service class instead of a single flow, it is important to model and characterise the external and the internal (or internode) traffic flows on a *per class* basis in order to plan and manage these networks to meet performance objectives (QoS) as required by the various traffic classes.