



Centrum voor Wiskunde en Informatica

REPORTRAPPORT

PNA

Probability, Networks and Algorithms



Probability, Networks and Algorithms

Bandwidth trading under misaligned objectives:
decentralized measurement-based control

M.R.H. Mandjes, M. Ramakrishnan

REPORT PNA-E0614 DECEMBER 2006

Centrum voor Wiskunde en Informatica (CWI) is the national research institute for Mathematics and Computer Science. It is sponsored by the Netherlands Organisation for Scientific Research (NWO). CWI is a founding member of ERCIM, the European Research Consortium for Informatics and Mathematics.

CWI's research has a theme-oriented structure and is grouped into four clusters. Listed below are the names of the clusters and in parentheses their acronyms.

Probability, Networks and Algorithms (PNA)

Software Engineering (SEN)

Modelling, Analysis and Simulation (MAS)

Information Systems (INS)

Copyright © 2006, Stichting Centrum voor Wiskunde en Informatica
P.O. Box 94079, 1090 GB Amsterdam (NL)
Kruislaan 413, 1098 SJ Amsterdam (NL)
Telephone +31 20 592 9333
Telefax +31 20 592 4199

ISSN 1386-3711

Bandwidth trading under misaligned objectives: decentralized measurement-based control

ABSTRACT

This paper studies the interplay between a profit-maximizing network and a number of users competing for the finite bandwidth on each link. In our setting, the objectives of the network and the users are ‘misaligned’, in that the prices that optimize the network’s profit do not maximize the aggregate utility of the users. The links set the prices for bandwidth and the users react to them by revealing their preferred amounts of bandwidth. A first contribution of this work is an iterative procedure for a single-link model. In this provably convergent scheme, the link adapts the price to achieve profit maximization, and each individual user adapts its demand for bandwidth so as to maximize its ‘compensated utility’, where utility is a function of its allocated bandwidth. Importantly, the scheme relies on communication between the link and the individual users, but *not* between users. In practice, the utility of the users (ISP s) strongly depends on the level of satisfaction experienced by their clients (the end-users). We show how the iteration scheme can be adapted to the more natural situation of utility being a function of the loss probability, rather than a function of the bandwidth. Since the end-users’ supply of traffic is not fully known to the ISP s, we develop a Bayesian approach for estimating the loss probability from measurements; we do so in the practically relevant context of Gaussian input traffic. The resulting estimator proves to be particularly useful for risk-averse ISP s.

2000 Mathematics Subject Classification: 60K25

Keywords and Phrases: Profit maximization, bandwidth trading, quality of service, measurement-based control, Gaussian queues

Note: Part of this work was carried out when M. Mandjes was visiting the University of Melbourne

Bandwidth trading under misaligned objectives: decentralized measurement-based control

M. Mandjes* and M. Ramakrishnan†

December 13, 2006

Abstract

This paper studies the interplay between a profit-maximizing network and a number of users competing for the finite bandwidth on each link. In our setting, the objectives of the network and the users are ‘misaligned’, in that the prices that optimize the network’s profit do not maximize the aggregate utility of the users. The links set the prices for bandwidth and the users react to them by revealing their preferred amounts of bandwidth.

A first contribution of this work is an iterative procedure for a single-link model. In this provably convergent scheme, the link adapts the price to achieve profit maximization, and each individual user adapts its demand for bandwidth so as to maximize its ‘compensated utility’, where utility is a function of its allocated bandwidth. Importantly, the scheme relies on communication between the link and the individual users, but *not* between users.

In practice, the utility of the users (ISP s) strongly depends on the level of satisfaction experienced by their clients (the end-users). We show how the iteration scheme can be adapted to the more natural situation of utility being a function of the loss probability, rather than a function of the bandwidth. Since the end-users’ supply of traffic is not fully known to the ISP s, we develop a Bayesian approach for estimating the loss probability from measurements; we do so in the practically relevant context of Gaussian input traffic. The resulting estimator proves to be particularly useful for risk-averse ISP s.

Keywords: Profit maximization, bandwidth trading, quality of service, measurement-based control, Gaussian queues

*M. Mandjes (email: mmandjes@science.uva.nl) is with Korteweg-de Vries Institute, University of Amsterdam, Plantage Muidergracht 24, 1018 TV Amsterdam, the Netherlands. MM is also affiliated to CWI, Amsterdam, the Netherlands, and EURANDOM, Eindhoven, the Netherlands.

†Corresponding Author: M. Ramakrishnan (email: m.ramakrishnan@ms.unimelb.edu.au) is with the ARC Centre of Excellence for Mathematics & Statistics of Complex Systems, The University of Melbourne, Parkville, VIC, 3010, Australia. +613 8344 6797

1 Introduction

For reasons of scalability, many on-line systems are organized in a decentralized way. It is clearly of much interest to devise procedures that help improve their performance, but at the same time are relatively simple to implement. Importantly, in a decentralized setting, control procedures should be based on ‘local information’ only. Thus a rather restrictive constraint is imposed on the set of feasible control strategies.

In the context of modern communications networks several examples of on-line systems with decentralized control can be found. Consider for instance a (broadband) network that sells resources (bandwidth in the setting of the present paper) to Internet Service Providers (ISP s; in this paper ‘users’). These ISP s, in turn, sell Internet connectivity to clients in the corporate and residential market (in this paper ‘end-users’). The ISP s need network bandwidth in order to be able to offer the end-users an acceptable Quality of Service (QoS), as agreed upon in a Service Level Agreement (SLA); for a further discussion of this hierarchy, see [4]. The amount of bandwidth each ISP chooses to buy should maximize its *utility*, where utility expresses the extent to which the end-users can be guaranteed their SLA, compensated by the price to be paid by the ISP to the network. However, the network has a potentially conflicting objective. When selling resources to the ISP s, it usually aims to maximize *profit*. Since the amount of bandwidth on each link of the network is finite, the ISP s are essentially competing for this scarce resource.

The question that naturally arises is that of the coordination of the competition between the users. A complicating issue is that in practice the users cannot communicate directly with each other. Instead, they rely on individual communication with the network, for each of the links that they traverse. More precisely, the network sets the prices (which they charge the users per supplied unit of bandwidth), and each user reacts to these by revealing their preferred amount of bandwidth. This is done without knowing the reaction of the other users. The network can then adapt the prices and the process continues. Clearly, it is desirable that these (iterative) price and bandwidth adaptations are such that convergence is guaranteed to the utility-maximizing allocation and profit-maximizing price. The development of such a procedure is a main contribution of this paper.

The bandwidth allocation reached reflects both the users’ and the network’s objectives. One could term these objectives ‘aligned’, if it is in the network’s interest to set prices that ensure that the equilibrium allocation also maximizes the users’ aggregate utility; this is often referred to as ‘social welfare optimization’. An example of such an aligned setting is studied by Kelly [10, Section 2]; its convergence to an equilibrium solution (by applying an iterative optimization-based flow control scheme) was analyzed in detail by Low and Lapsley [13], and Kelly *et al.* [11], whereas Johari and Tan [9] and Vinnicombe [21] addressed related stability issues. In fact, the setup of [10] enforces that the users and the network coordinate their actions so that the social optimum is reached. One can

alternatively consider the situation in which the network imposes a utility function on each route that achieves a certain fairness criterion [16].

It is not realistic to assume that the network and the users always have aligned objectives. As explained above, this paper focuses on such a ‘misaligned situation’, in which the network acts as a monopolist and wishes only to maximize its profit (subject to the link capacity constraints), while each individual user optimizes its *compensated* utility. We study a model in which the network acts as a price-setter, while the users act as price-takers, and adjust their desired bandwidths according to the prevailing price. In this scenario, both parties’ objectives have an impact on the final allocation of capacity, and corresponding prices. It is therefore natural to expect that the resulting prices and allocations will deviate from the ones under social welfare maximization.

The differences between the outcomes of social welfare optimization and monopoly markets have been well researched in the economic literature, see in particular the seminal work by Pigou [18]. Informally, one could say that monopoly markets are usually beneficial for the supplying party (the network, in our situation), but create a ‘welfare loss’ for the participants (the users, in our situation). Recently, the term *price of anarchy* was introduced to quantify this welfare loss [19]; it reflects the inefficiency of the equilibrium solution. Huang, Ozdağlar and Acemoğlu [8] showed, in a given setting, that for specific types of utility functions, the monopoly equilibrium coincides with the social welfare equilibrium, with regard to bandwidth allocations. Ozdağlar [17] considered the price of anarchy in the context of routing and flow control in a network consisting of multiple parallel links, where a cost per unit bandwidth and a congestion-based cost were imposed on each user. For a utility curve satisfying stated conditions, she derived a bound on the efficiency of the solution by comparing the equilibria reached in the monopolist and social welfare optimization contexts, respectively.

Instead of considering the price of anarchy, we focus our attention on devising an iteration scheme by which the network and users achieve their respective equilibrium solutions. We impose the condition that the solution be found in a distributed manner; importantly, each user does not have any information about the other users’ preferences. We derive an optimization-based scheme for price updates (network) and bandwidth updates (users), which is provably convergent. In the first part of this paper it is assumed that the users’ utilities are functions of the bandwidth allocated.

As discussed, the users (ISP’s) have guaranteed their clients (end-users) a certain performance, for instance in terms of a packet loss probability. For this reason it may be more natural to consider a setup in which the user’s utility curve is not a function of the bandwidth x , but rather of the loss probability $\delta \equiv \delta(x)$ that is realized when the user is allocated x . To be able to use such a framework, the users must have a procedure to estimate the loss probability $\delta(x)$ as a function of the allocation x . However, $\delta(x)$ depends

not only on x , but also on the characteristics of the traffic offered by the end-users; for a given x , the loss probability is larger when the end-users generate higher loads, or when the traffic rate fluctuates more fiercely in time.

The reasoning of the previous paragraph motivates the need for the ISPs to measure the traffic of their clients, so as to estimate the function $\delta(x)$ on-line, and then to insert this estimate in their (loss-probability-based) utility function. We propose a procedure that estimates $\delta(x)$ in a Bayesian way [6], which yields conservative estimates, and is therefore particularly useful for risk-averse ISPs. The analysis is based on the assumption of Gaussian traffic; this is justified as long as the traffic aggregation level (for instance in terms of numbers of end-users) is sufficiently high [3, 12, 15], which is typically the case for large ISPs.

The organization of this paper is as follows. In Section 2 we present our model and some preliminaries. In Section 3 we analyze the situation in which the utility curves are functions of the bandwidth allocated, and present a provably convergent iterative scheme for the prices and the allocations. Section 4 adapts this iterative scheme to the situation in which each ISP's utility is a function of the loss probability; it is explained in detail in Section 5 how the loss probability may be estimated from measurements. In Section 6 the procedures are illustrated through numerical experiments. Section 7 presents concluding remarks.

2 Model and preliminaries

In this section, we present our model of a profit-maximizing network that is used by utility-maximizing routes. The problem is stated in microeconomic terms, and a number of preliminaries and first observations are given.

Links. Let \mathcal{L} denote the set of all physical links and assume each link $\ell \in \mathcal{L}$ has a finite service capacity (i.e., bandwidth) c_ℓ . Each link $\ell \in \mathcal{L}$ has the authority to set a price p_ℓ for each unit of bandwidth, with the aim of maximizing the total network profit.

Routes. A route is defined as a subset of the links, starting at an origin node and ending at a destination node. Let \mathcal{R} denote the set of routes in the network. Each route $r \in \mathcal{R}$ is allocated the same amount of bandwidth at any link that it traverses, and when allocated x_r units, it derives a utility $U_r(x_r)$. Routes can be identified with users (ISPs). The routes are 'at the mercy' of the network, in the sense that they request capacity given the prevailing prices; they choose the amount of capacity that maximizes their compensated utility (defined as utility less cost; a precise definition follows later). We assume that the routes cannot co-operate; in this way, we preclude the possibility that routes can collude to try and bring the prices down.

Identifying routes with ISPs, the utility optimization described above reflects the route's task to satisfy the Service Level Agreement (SLA) that it has agreed upon with the end-users. Clearly, the extent to which the ISP is able to meet the SLA of its customers decreases as the loss probability increases, whereas the loss fraction decreases in the amount of bandwidth allocated. Therefore it is natural to assume that, for every ISP r , the utility curve $U_r(x_r)$ increases in x_r .

As is customary, cf. [10], we assume that these utility functions are strictly increasing and concave in x_r . In addition, for reasons that will become clear below, we assume that the utility functions are differentiable, with derivatives that are one-to-one.

If the network sets prices that are too high, then the routes will request very low amounts of bandwidth. The profit derived from this allocation will not be large, and the network is not satisfied. Conversely, if the prices are very low, the capacity requested will be larger, however possibly not large enough to offset the decreased prices. In addition, the routes' requests may exceed the finite link capacity. Therefore, there may be an optimal set of prices that balance these opposing forces. Importantly, observe that the objectives of the various ISPs are not aligned with the network's objective. Let us now look in more detail at a formal description of these objectives.

User optimization problem. As stated above, routes (or users) occupy a subset of the physical links. As the route is meant to transport a stream of data packets, it requires the same amount of bandwidth from each link that it traverses. A route $r \in \mathcal{R}$ is faced with the total price of $x_r \sum_{\ell \in r} p_\ell$ when acquiring x_r units of bandwidth. Since we assume that the routes cannot co-operate with each other, they make their decision solely based on their own utility function and the prevailing prices. Hence the *user optimization problem* is given by

$$\max_{x_r \geq 0} U_r(x_r) - x_r \cdot \sum_{\ell \in r} p_\ell. \quad (1)$$

Let $U'_r(x_r)$ denote the derivative of the utility function for route r with respect to x_r , and $V_r(x_r)$ denote the inverse of this derivative. By the conditions imposed on the utility functions, the solution to the user problem can be expressed as (with $p^r := \sum_{\ell \in r} p_\ell$)

$$x_r(p^r) := V_r(p^r), \quad (2)$$

assuming the right-hand side of (2) is non-negative. In essence, $x_r(p^r)$ defines a *demand function* for bandwidth for each route r , which, as expected, decreasing in price. We make the added assumption that the utility functions induce demand functions that are twice continuously differentiable with respect to p^r .

Network optimization problem. Now we know how user r reacts to price p^r , the next question is how the prices are set. Recall, the links must set p_ℓ to maximize its profit in a

distributed manner. If the routes request capacity equal to $\vec{x} \equiv (x_r, r \in \mathcal{R})$, then the profit equals the sum of the revenues earned on each link: $\sum_{\ell \in \mathcal{L}} (p_\ell \cdot \sum_{r: \ell \in r} x_r)$; for ease we leave out the cost to the network in providing bandwidth to the routes. Since the capacities \vec{x} are defined by the demand curve (2), the profit can be expressed as a function of the price, with $\vec{p} \equiv (p_\ell, \ell \in \mathcal{L})$:

$$\pi(\vec{p}) = \sum_{\ell \in \mathcal{L}} \left(p_\ell \cdot \sum_{r: \ell \in r} x_r(p^r) \right). \quad (3)$$

While the links choose prices that maximize their objective (3), they must also adhere to a number of constraints. In the first place there must be constraints requiring that the capacity requested by the routes on a specific link is no larger than the link capacity. In addition, we impose a cap on the the prices that each link may charge, p_{\max} ; for ease we assume that this maximum price is the same for each link. Such a maximum price may be due to regulatory reasons, but we illustrate in Section 3 that it is useful from a mathematical perspective as well. We denote the vector of price caps p_{\max} . We thus obtain the following *network optimization problem*

$$\begin{aligned} & \max_{\vec{p}} && \pi(\vec{p}) \\ & \text{subject to} && \sum_{r: \ell \in r} x_r(p^r) \leq c_\ell \quad \text{for } \ell \in \mathcal{L} \\ & && 0 \leq p_r \leq p_{\max} \quad \text{for } r \in \mathcal{R}. \end{aligned} \quad (4)$$

We assume that $\sum_{r: \ell \in r} x_r(\sum_{\ell \in r} p_{\max}) < c_\ell$; this guarantees this if every link chooses its maximum price, the resulting capacity requests will be feasible.

The formulation (4) is a static optimization problem. So from the network's perspective, if it could be solved centrally, the profit-maximizing prices could be set, and the problem solved. However, there are issues that preclude this from happening.

- First, we seek a distributed implementation. We assume that there is no central controller of the network, and hence any decision regarding the change in prices must be performed in a manner that does not require direct communication between the links.
- In addition, observe that the formulation given in (4) combines the prices charged by the links and the demand functions of the routes. We assume that the demand functions are not known by the links; they can only 'ask' the routes to reveal how much bandwidth they want at posted prices. This implies that any solution of this problem requires iterative feedback between the links and the routes that traverse it.

The idea of reaching the solution dynamically is even more pertinent when the routes experience a change in their utility functions (for example because the traffic

offered by the end-users changes; when the end-users' load increases more capacity is needed to meet the SLAs, and therefore the utility assigned to a given capacity decreases).

The intuitive explanation of a distributed mechanism is that each link should learn the demand curves of the routes during some sort of bidding process: the network advertises a set of prices \vec{p}_1 , and sees the responses $x_r(\vec{p}_1)$ for $r \in \mathcal{R}$; then the network changes the prices to \vec{p}_2 , and sees the new responses $x_r(\vec{p}_2)$ for $r \in \mathcal{R}$, etc. This bidding process must be automated, that is, performed through agents, to be practical in a network setting. Importantly, the algorithm that will be presented in the next section allows a distributed implementation.

3 Single-link network

In this section, we consider the network as a single link of capacity C , shared by a set of ISPs, each denoted $r \in \mathcal{R}$. For this single-link model, we devise a distributed algorithm. The reason for commencing with a simplified network is twofold. First, analyzing the single-link model allows immediate insight into the differences between social welfare maximization and the profit maximization scenario that we explore in this paper. These insights are detailed in Section 3.1. Second, conditions required for convergence of the devised distributed algorithm are simpler to check in the single-link case. We make a remark regarding extension to the network setting at the conclusion of this section.

Because we concentrate on a single link, we may suppress the dependence on ℓ , and we denote by p the price of bandwidth. Thus the optimization problem user r is faced with is $\max_{x_r} U_r(x_r) - x_r p$, and the demand curve is given by $x_r(p) = V_r(p)$. For the sake of completeness, we translate the formulation (4) to the single link case. The network problem is given by

$$\begin{aligned} \max_p \quad & p \cdot \sum_{r \in \mathcal{R}} x_r(p) \\ \text{subject to} \quad & \sum_{r \in \mathcal{R}} x_r(p) \leq C \\ & 0 \leq p \leq p_{\max}. \end{aligned} \tag{5}$$

Since each demand function is decreasing and continuous in p , the aggregate demand $f(p) := \sum_r x_r(p)$ is also. Therefore, the link capacity constraint defines a minimum price p_{\min} below which the constraint is violated. Of course, p_{\min} could be negative. If this is the case, the non-negativity constraint makes the link capacity constraint redundant. We assume that $p_{\min} \neq 0$ for technical reasons.

Since the feasible region is closed and $\pi(p)$ is continuous over it, there exists an optimal solution. If the restriction that $p \leq p_{\max}$ was omitted, it is possible that no finite optimum exists. Take for example the case where $U_r(x_r)$ is concave, but $x_r(p)$ is convex with

$x_r(p) \rightarrow \kappa$, a finite value, when $p \rightarrow \infty$. Then the link can make the price arbitrarily high and $\pi(p) \rightarrow \infty$. This shows that the introduction of p_{\max} serves a mathematical purpose, besides reflecting the regulatory environment.

3.1 System Optimization versus Profit Optimization

Before describing the iterative scheme, we first consider, through two insightful lemmas, the differences between the equilibrium solutions of the system-optimizing (social-welfare) and profit-maximizing problems in the single-link case, respectively.

The optimization-based flow control that Kelly *et al.* [11] and Low and Lapsley [13] performed are examples of system optimization, where the network acts in such a way that the equilibrium maximizes aggregate utility, $\sum_r U_r(x_r)$. However, the equilibrium vector of bandwidths $\vec{x}^{(s)}$ resultant from these scenarios may not be the same as that attained via the profit-maximizing problem described above, denoted $\vec{x}^{(\pi)}$. Under system optimization *all* of the link capacity C is allocated to the routes, under profit optimization this is not necessarily the case.

Lemma 3.1 *The equilibrium reached yields an aggregate utility*

$$\sum_r U_r(x_r^{(s)}) \geq \sum_r U_r(x_r^{(\pi)}),$$

where this inequality can be strict. In economic terms, the profit maximization scenario creates a possible welfare loss because the network acts as a monopolist.

Proof. Consider the objective function $\pi(p)$ of the link problem (5). Let $x'_r(p)$ denote the derivative of route r 's demand function with respect to p . Suppose that there exists a maximizer $0 < p^* < p_{\max}$, where p^* satisfies

$$\pi'(p^*) = p^* \sum_{r \in \mathcal{R}} x'_r(p^*) + \sum_{r \in \mathcal{R}} x_r(p^*) = 0.$$

Since this is independent of the total available bandwidth C , the link capacity constraint may or may not be active at the optimum and therefore there may be bandwidth on the link left unallocated at p^* . If spare capacity exists, aggregate utility is strictly less than in the welfare optimization case, since aggregate utility could be increased by allocating the excess capacity to one of the ISP s.

We analyze separately the cases in which the utility functions do not admit a solution to the above equation. If the cap on prices p_{\max} comes into effect, then by assumption $\sum_r x_r(p_{\max}) < C$, and hence there is unallocated link capacity at equilibrium. If $p^* = 0$, then the link capacity constraint is inactive by assumption. Hence, again, there is unallocated capacity at equilibrium. Therefore, by the same argument as above, aggregate utility is strictly less than in the welfare-maximizing scenario.

Finally, if $p^* = p_{\min}$, then all link capacity is allocated to the ISP s. The aggregate utility can only be as great as $\sum_r U_r(x_r^{(s)})$, but no larger. \square

Lemma 3.2 *The allocated bandwidth to each route under profit maximization is not greater than the allocated bandwidth in the system (social-welfare) optimization case. That is,*

$$x_r^{(\pi)} \leq x_r^{(s)} \quad \text{for all } r \in \mathcal{R}.$$

Proof. Consider the optimal allocations under the profit-maximizing and the aggregate utility-maximizing scenarios; $x_r^{(\pi)}$ and $x_r^{(s)}$, respectively. The latter can be mapped to some price q for which $U'_r(x_r^{(s)}) = q$. The maximum profit to the link is

$$\pi(p^*) = p^* \sum_{r \in \mathcal{R}} x_r^{(\pi)} \geq q \sum_{r \in \mathcal{R}} x_r^{(s)},$$

since the left-hand side is optimal.

Since the total allocated bandwidth in the profit scenario is not greater than for social optimization, we must have that $p^* \geq q$. Since each demand function is a decreasing function of price, $x_r^{(\pi)} = x_r(p^*) \leq x_r(q) = x_r^{(s)}$. \square

These results are intuitive; in economic terms, they highlight the difference between a monopoly market and perfect competition, and the welfare loss that results from the former.

Corollary 3.1 It follows from Lemma 3.2 that if all link capacity is allocated, then it is done so in a system (socially) optimum way. \diamond

3.2 Existence of a solution

As the single-link problem is one of constrained optimization with respect to a one-dimensional variable p , the constraints reduce to a feasible interval from $\max\{0, p_{\min}\}$ to p_{\max} . Since $p_{\min} \neq 0$, at any time at most one constraint is active. Therefore, trivially, there is no linear dependence between the constraint gradients of the active constraints and hence any local maximizer is regular [1, Section 3.3].

The concavity of the utility functions is not sufficient to guarantee a unique optimal solution of (5). Recall, the objective function and constraints of the link problem are written in terms of the demand functions $x_r(p)$ of each route. The assumptions on the utility functions guarantee that the demand functions $x_r(p)$ exists, but do not imply that they are concave (or convex). With $\pi(p) = pf(p)$, we have that

$$\pi''(p) = 2f'(p) + pf''(p);$$

$f'(p)$ is negative, but $f''(p)$ could be both positive and negative. We conclude that $\pi(p)$ is not necessarily concave, and hence we cannot guarantee the existence of a unique optimum.

Let $\mathcal{L}(p, \eta, \gamma)$ denote the Lagrangian of the network problem, with η the multiplier of the link capacity inequality constraint and γ the vector of multipliers of the price inequality constraints:

$$\mathcal{L}(p, \eta, \gamma) := pf(p) - \eta(f(p) - c) - \gamma_1(p - p_{\max}) + \gamma_2 p. \quad (6)$$

Recall p_{\min} was defined as the solution to $f(p) = c$ and hence is implicit in the statement of the Lagrangian.

We proceed by assuming the *second-order sufficient conditions* (SOSC) for inequality constrained problems hold [1, Proposition 3.3.2], which can be described as follows. Let ∇_p denote the derivative operator with respect to p . Then the SOSC state that the triple (p^*, η^*, γ^*) describes a local maximum if it satisfies

$$\nabla_p \mathcal{L}(p^*, \eta^*, \gamma^*) = 0, \quad f(p^*) \leq c, \quad 0 \leq p^* \leq p_{\max}, \quad \eta^* \geq 0, \quad \gamma^* \geq 0;$$

where

$$\eta^* = 0, \quad \gamma^* = 0 \quad \text{if the corresponding constraint is inactive;}$$

$$\eta^* > 0, \quad \gamma^* > 0 \quad \text{if the corresponding constraint is active ('strict complementarity');}$$

$$y^T \nabla^2 \mathcal{L}(p^*, \eta^*, \gamma^*) y < 0,$$

for all vectors $y \neq 0$ such that the active-constraint gradients at p^* are perpendicular to y .

3.3 Iterative scheme

As discussed earlier, distributed implementation requires that the optimum of (5) is found without the link having complete knowledge of each ISP's demand function, and each ISP being unaware of the demand functions of the other ISP's. Instead, updates to both the prices and the allocated bandwidths must be performed iteratively, as reactions to one another. We require that the information transferred between the ISP and the link be local.

Since there is a separation between the users and the link, a first-order method is more appropriate; second-order methods (such as Newton-type schemes) typically involve non-local information. In the general network setting, first-order methods also fit with the requirement that links do not communicate directly with one another. The most basic first-order method involves the Lagrangian defined in Section 3.2. Its convergence, however, relies on the Hessian being concave at the local maximum. This is stronger than the second-order sufficient conditions and hence may not be satisfied.

For the single-link network, we proceed by defining the *augmented* Lagrangian

$$\mathcal{L}_m(p, \eta) := \pi(p) - \frac{1}{2m} \left((\max\{0, \eta + m(f(p) - c)\})^2 - \eta^2 \right),$$

where only the link capacity constraint is involved. The form of the second term (a penalty function) is derived by appending slack variables to the inequality constraints, adding a penalty for violation and finally maximizing over the slack variables analytically [1, Section 4.2].

Assume the SOSC hold. Then by [14, Ch. 13, Prop. 1], there exists an \bar{m} such that for all $m > \bar{m}$, the augmented Lagrangian $\mathcal{L}_m(p^*, \eta^*)$ has a local maximum at p^* with penalty parameter m . First-order Lagrange methods, while normally described using the original Lagrangian (6), can be applied with the augmented version also, when $m > \bar{m}$. The augmented Lagrangian not only allows us to use first-order methods, but imposes a penalty m on infeasibility. This has the benefit of discouraging iterates to stray too far from the feasible region.

Let p_j denote the price per unit bandwidth at the j -th iteration, and η_j denote the multiplier value at the j -th iteration. The first-order augmented Lagrangian method is defined through the updates

$$\eta_{j+1} := \max \left\{ 0, \eta_j + \epsilon (f_{j+1} - c) \right\}, \quad (7)$$

$$p_{j+1} := \left[p_j + \epsilon \left(f_{j+1} + f'_{j+1} [p_j - \max\{0, \eta_j + m(f_{j+1} - c)\}] \right) \right]_0^{p_{\max}}, \quad (8)$$

where $f_{j+1} := \sum_r x_r(p_j)$ and $f'_{j+1} := \sum_r x'_r(p_j)$; also, $[x]_b^a$ is defined as $\min\{\max\{x, a\}, b\}$. There exists a neighborhood around (p^*, η^*) and an $\bar{\epsilon}$ such that for all $\epsilon \in (0, \bar{\epsilon}]$, (p^*, η^*) is a point of attraction of the updates given by (7)-(8). This follows from a variation of Proposition 4.4.2 in [1]. Details of the proof of convergence for a first-order method applied to the augmented Lagrangian for inequality-constrained problems can be found in the Appendix.

The update formulae (7)-(8) involve the demand functions for all routes $r \in \mathcal{R}$, through the f_{j+1} and f'_{j+1} terms. However, as discussed earlier, there is a separation between the link and the ISPs. Hence, the network does not know the individual ISP's utility functions (and hence demand curves). Therefore the function $f(p)$ must be learned by the link at each iteration, by combining the requests that each route submits. The price update p_{j+1} not only involves f_{j+1} , the current total demand for bandwidth at the prevailing price, but also its derivative f'_{j+1} . This could be implemented by the requirement that each ISP, in addition, submits an 'elasticity' $x'_r(p_j)$ at each iteration. Since each ISP knows its own utility function, this is not an unrealistic requirement. The algorithm for the single link network could operate as shown below.

Algorithm 3.1 — Profit Maximization

1. At iteration j , the link sets its price p_j and broadcasts it to the routes that traverse it.
2. Each ISP $r \in \mathcal{R}$ sees the current price and responds with two pieces of information:

$$x_r(p_j) = V_r(p_j) \quad \text{and} \quad x'_r(p_j) = \frac{d}{dp} V_r(p_j).$$

3. Through agents, these requests are aggregated and the link sees a total desired bandwidth of f_{j+1} and its corresponding derivative f'_{j+1} .
4. Using the updates (7)-(8), the link updates its price to p_{j+1} and its multiplier to η_{j+1} .
5. It re-broadcasts a price of p_{j+1} , and the process continues until a pre-specified convergence criterion is met.

Remark 3.1 In the setting described in this (and the previous) section, it is assumed that the utility functions $U_r(x_r)$ are increasing and convex for $x_r \geq 0$. It can be questioned, however, whether this is a realistic assumption in practical situations. In fact each ISP has to be capable of handling a certain average rate that is generated by the end-users, say μ_r on route r ; if the ISP is allocated less than μ_r , then the ISP fails to transmit all traffic, and therefore it is unlikely that SLA s are met. Hence each route r requires at least bandwidth μ_r , and on top of that some extra bandwidth x_r can be allocated. In other words: the total bandwidth allocated, say \bar{x}_r , is the sum of the average rate μ_r (needed to ensure that all incoming traffic can be handled) and ‘excess bandwidth’ x_r (needed to provide QoS; the higher x_r , the better the QoS delivered). Therefore it is natural to assume that in this context the utility function is strictly increasing and concave in the *excess* bandwidth x_r (rather than the total bandwidth \bar{x}_r).

It is easy to adapt the Algorithm 3.1 to the situation described above. The user maximizes the utility of the excess bandwidth x_r less the cost of the bandwidth $\bar{x}_r = \mu_r + x_r$,

$$\max_{x_r \geq 0} U_r(x_r) - (\mu_r + x_r)p,$$

whereas the network solves

$$\max_p p \cdot \sum_{r \in \mathcal{R}} (\mu_r + x_r(p)) \quad \text{under} \quad \sum_{r \in \mathcal{R}} (\mu_r + x_r(p)) \leq C,$$

and $p \in [0, p_{\max}]$; this yields the optimal price p^* , so that the total assigned bandwidth to user r is $\mu_r + x_r(p^*) = \bar{x}_r$.

Using the same approach as for the original problem, we should use the updates

$$\eta_{j+1} := \max \left\{ 0, \eta_j + \epsilon \left(f_{j+1} + \bar{\mu} - C \right) \right\}, \quad (9)$$

$$p_{j+1} := \left[p_j + \epsilon \left(f_{j+1} + \bar{\mu} + f'_{j+1} [p_j - \max\{0, \eta_j + m(f_{j+1} + \bar{\mu} - C)\}] \right) \right]_0^{p_{\max}}, \quad (10)$$

where f_{j+1} and f'_{j+1} are the total demand and derivative of total demand for *excess* bandwidth, respectively, and $\bar{\mu} := \sum_{r \in \mathcal{R}} \mu_r$. \diamond

Remark 3.2 The first-order method can be extended to the general network setting; then a vector of prices should be updated. Convergence to a local maximum \bar{p}^* can be guaranteed if the second-order sufficient conditions are satisfied for the larger problem, and it can be shown that \bar{p}^* is regular. As expected, this is more complicated in the network case. We do not include further details here. \diamond

In this section, we have assumed that each ISP assigns utility to bandwidth. As argued in the introduction, it may be more natural to consider the situation that utility is a function of the loss probability, as this is what the ISP has negotiated with its end-users. The next section addresses the optimization problem for this situation.

4 Gaussian traffic; utility as a function of loss

In this section we briefly review the concept of Gaussian queues, and provide some examples of possible utility functions. These examples show that the utility curve can be considered a function of the loss probability δ , rather than the bandwidth x_r allocated to route r . Finally we explain how to change Algorithm 3.1 so that it can be used in this situation.

4.1 Gaussian queues

As argued by [3, 12, 15], aggregate traffic streams in packet networks can be accurately approximated by Gaussian processes, in particular when the traffic aggregation level (for instance in terms of numbers of users) is sufficiently high. This class of models covers a broad range of correlation structures, including long-range dependence.

The total amount of traffic $A_r(t)$ offered to the r -th ISP in an interval $[0, t]$ is Normally distributed with mean $\mu_r t$ and variance $v_r(t)$. Suppose this ISP is allocated an amount of bandwidth $\bar{x}_r = \mu_r + x_r$ (cf. Remark 3.1). With the assumption of Gaussianity, we focus on the performance measure

$$\delta(x_r) := \mathbb{P}(A_r(T) \geq (\mu_r + x_r)T), \quad (11)$$

that is, the probability that the total offered traffic over this chosen time-scale exceeds the allocated bandwidth. The choice of an appropriate T is primarily a task of the network manager, and depends on the types of applications being used; for further guidelines, see [20]. We rescale time such that $T \equiv 1$; for brevity we write $v_r \equiv v_r(1)$. The quantity $\delta(x_r)$ can be approximated as in [20]: the Chernoff bound yields

$$\delta(x_r) \approx \exp\left(-\frac{x_r^2}{2v_r}\right), \quad (12)$$

for $x_r > 0$, and 1 else.

As argued in Remark 3.1, one has to allocate to route r minimally the mean rate μ_r , and in addition an excess rate x_r that can be determined by the iterative algorithm of the previous section. In practice, however, ISP's will assign utility to the loss probability that they experience rather than to (excess) bandwidth as such; recall that this is because it is

the loss probability that they agreed upon with their clients (in the SLA). We can define, for every function $U_r(\cdot)$, a function $\bar{U}_r(\cdot)$ such that

$$U_r(x_r) = \bar{U}_r(\delta(x_r)) = \bar{U}_r\left(\exp\left(-\frac{x_r^2}{2v_r}\right)\right);$$

the last equality is based on Approximation (12).

Example 4.1 In this example we show that natural choices for $\bar{U}_r(\cdot)$ lead to functions $U_r(\cdot)$ that satisfy the assumptions made in Section 2. Consider the following choices:

(1) $\bar{U}_r(\delta) = \alpha(-\log \delta)^\beta$, for positive α and $\beta \in (0, \frac{1}{2})$. Then

$$U_r(x_r) = \frac{\alpha}{(2v_r)^\beta} x_r^{2\beta},$$

which is indeed strictly increasing and concave in x_r .

(2) $\bar{U}_r(\delta) = \alpha - \beta/(1 - \delta)$ for positive β . Now

$$U_r(x_r) = \alpha - \beta \left(1 - \exp\left(-\frac{x_r^2}{2v_r}\right)\right)^{-1},$$

also strictly increasing and concave. ◇

We observe that we can rewrite the user optimization problem, for the single link, to $\max_{x_r \geq 0} \bar{U}_r(\delta(x_r)) - (x_r + \mu_r)p$. To perform this optimization, we should adapt Algorithm 3.1 slightly, as will be explained in the next section.

4.2 Characterization of the demand curves

In Algorithm 3.1 the routes' utilities depend on the bandwidth that they are allocated. We now describe how the algorithm should be adapted to the situation where utility curves $\bar{U}_r(\cdot)$ depend on the loss probability.

Recall, Step 2 of the algorithm required each ISP to tell the link (through agents) its current demand at the prevailing price and its current derivative of demand at the prevailing price. We now discuss how the ISP s can use knowledge of the function $\delta(x_r)$ to calculate these quantities. As we have seen earlier, the demand curve is defined as the solution to the equation $U_r'(x_r) = p$. By assumption, we know a solution exists (see Section 2). Since $U_r(x_r) = \bar{U}_r(\delta(x_r))$, by the chain rule, we have

$$U_r'(x) = \bar{U}_r'(\delta(x_r)) \frac{d}{dx_r} \delta(x_r). \tag{13}$$

By equating Expression (13) to p , we find the demand for excess bandwidth at price p (denoted $x_r(p)$); in the algorithm, it can be found numerically, for example, via a Golden Section search. The algorithm also requires $x_r'(p)$; we can approximate this by

$$\frac{d}{dp} x_r(p) \approx \frac{x_r(p+h) - x_r(p)}{h}, \tag{14}$$

for a small $h > 0$. These quantities can then be used in Step 2 of the algorithm, with updates (9)-(10) used to adapt the price.

The above approach clearly requires that the ISP *know* their required mean rate μ_r and the function $\delta(x_r)$. In a practical setting, however, the ISP *s* do not *a priori* know the μ_r and v_r of their customers. Therefore they may rely on measurements to estimate this loss probability. In the next section we present two ways to estimate $\delta(x_r)$ from measurements.

5 Estimation of loss probabilities

In this section we describe two approaches to estimate the loss probability from measurements. In the first place one could estimate the mean and variance, continue as if these are the true parameter values, and insert them in the approximation; this approach is called *certainty equivalence*. As will be discussed, it does not take into account the inherent randomness of the measurements, and therefore we propose a more conservative, Bayesian alternative, based on the so-called *Inverse Sanov* theorem. We conclude this section by showing how each ISP can use the Bayesian estimator to calculate the quantities needed in Step 2 of the algorithm in the context of Gaussian queues.

Consider a single ISP; it is offered an amount of traffic per unit time, say $A(1)$, that is Normally distributed with *unknown* mean μ and variance v . There are essentially two different ways to estimate $\Delta(x) := \mathbb{P}(A(1) > x)$. A ‘frequentist approach’ would be to first estimate the mean and variance of $A(1)$. This would be done in the classical way: measure traffic over non-overlapping windows of length 1, far enough apart to make sure that they are just weakly dependent. We thus obtain observations a_1, \dots, a_n . From this sample, we can compute the estimates $\hat{\mu}_n := n^{-1} \sum_{i=1}^n a_i$ and $\hat{v}_n := (n-1)^{-1} \sum_{i=1}^n (a_i - \hat{\mu}_n)^2$. Then one could proceed as if the estimates $\hat{\mu}_n$ and \hat{v}_n are the real parameter values and apply (12), which yields

$$\mathbb{P}(\mathcal{N}(\hat{\mu}_n, \hat{v}_n) \geq x) \approx \exp\left(-\frac{1}{2} \frac{(x - \hat{\mu}_n)^2}{\hat{v}_n}\right).$$

This idea leads to the following approximation.

Approximation 5.1 *The certainty-equivalence estimator of $\Delta(x)$ is*

$$\Delta_{\text{CE}}(x) := \exp\left(-\frac{1}{2} \frac{(x - \hat{\mu}_n)^2}{\hat{v}_n}\right) =: q(x \mid \hat{\mu}_n, \hat{v}_n). \quad (15)$$

We use the term ‘certainty-equivalence’ (CE), as it is based on the idea that the values $\hat{\mu}_n$ and \hat{v}_n are the true ones, thus ignoring the inherent uncertainty of the estimates. This explains why we expect that this approximation cannot be guaranteed to be conservative.

The Bayesian alternative would be the following. First realize that the event $\{A(1) > x\}$ can be considered as the combination of two effects. In the first place, it could be that the estimates $\hat{\mu}_n$ and \hat{v}_n of the mean and variance are too optimistic, i.e., lower than the true values μ_0 and v_0 . In the second place, it could be that, even if the estimates coincide with the true values μ_0 and v_0 , $A(1)$ is unusually high (a ‘large deviation’).

Whereas the second probability can be approximated by standard large-deviations techniques, the first probability can be estimated by applying the so-called *Inverse Sanov* theorem [6]. Suppose we know that the observations stem from a Normal distribution, and suppose their realizations reveal a sample mean of $\hat{\mu}_n$ and a sample variance of \hat{v}_n . We may be interested in the ‘probability’ $q_n(\mu_0, v_0 \mid \hat{\mu}_n, \hat{v}_n)$ that the observations actually stem from a Normal distribution with mean μ_0 and variance v_0 . One can show that the following approximation can be used:

$$q_n(\mu_0, v_0 \mid \hat{\mu}_n, \hat{v}_n) := \exp(-nI(\mu_0, v_0 \mid \hat{\mu}_n, \hat{v}_n)); \quad (16)$$

where

$$I(\mu_0, v_0 \mid \hat{\mu}_n, \hat{v}_n) := \frac{1}{2} \log \left(\frac{v_0}{\hat{v}_n} \right) + \frac{(\hat{\mu}_n - \mu_0)^2 + \hat{v}_n - v_0}{2v_0};$$

see [2, Thm. 6]. This statement, the Inverse-Sanov theorem, in fact says that the ‘probability’ of μ_0, v_0 deviating from $\hat{\mu}_n, \hat{v}_n$ vanishes exponentially in n . It is easily seen that $I = 0$ for $\mu_0 = \hat{\mu}_n$ and $v_0 = \hat{v}_n$ and positive elsewhere; I is a so-called ‘relative entropy’ or Kullback-Leibler distance.

We thus obtain the following first approximation, that ‘conditions’ on the values of μ_0 and v_0 . With q and q_n as defined in (15) and (16),

$$\Delta(x) \approx \int_{\mu_0, v_0} q(x \mid \mu_0, v_0) \cdot q_n(\mu_0, v_0 \mid \hat{\mu}_n, \hat{v}_n) d\mu_0 dv_0.$$

To further approximate $\Delta(x)$, we can apply the principle of the largest term (cf. [5, p. 25]): replace the integral by the maximum of the integrand. We thus obtain

$$\Delta(x) \approx \exp \left(-\frac{1}{2} Q_n(x, \hat{\mu}_n, \hat{v}_n) \right),$$

where

$$Q_n(x, \hat{\mu}_n, \hat{v}_n) := \min_{\mu_0, v_0} Q_n(x, \mu_0, v_0 \mid \hat{\mu}_n, \hat{v}_n);$$

$$Q_n(x, \mu_0, v_0 \mid \hat{\mu}_n, \hat{v}_n) := \frac{(x - \mu_0)^2}{v_0} + n \left(\log \left(\frac{v_0}{\hat{v}_n} \right) + \frac{(\hat{\mu}_n - \mu_0)^2 + \hat{v}_n - v_0}{v_0} \right).$$

The optimizing μ_0, v_0 can be found explicitly. The first order condition

$$\frac{d}{d\mu_0} Q_n(x, \mu_0, v_0 \mid \hat{\mu}_n, \hat{v}_n) = 2 \left(\frac{x + n\hat{\mu}_n - (n+1)\mu_0}{v_0} \right) = 0$$

immediately leads to

$$\mu_0^* = \hat{\mu}_n + \frac{x - \hat{\mu}_n}{n + 1},$$

whereas on the other hand

$$\frac{d}{dv_0} Q_n(x, \mu_0, v_0 \mid \hat{\mu}_n, \hat{v}_n) = -\frac{(x - \mu)^2}{v_0^2} + \frac{n}{v_0} - n \cdot \frac{(\hat{\mu}_n - \mu_0)^2 + \hat{v}_n}{v_0^2} = 0$$

leads to (after considerable calculus)

$$v_0^* = \hat{v}_n + \frac{(x - \hat{\mu}_n)^2}{n + 1}.$$

Inserting these into the objective function Q_n eventually yields the following approximation.

Approximation 5.2 *The Inverse-Sanov estimator of $\Delta(x)$ is*

$$\Delta_{\text{IS}}(x) := \left(1 + \frac{1}{n + 1} \cdot \frac{(x - \hat{\mu}_n)^2}{\hat{v}_n} \right)^{-n/2}. \quad (17)$$

It is interesting to see how this estimate behaves as a function of n . In the first place, we observe that for small n , the estimate μ_0^* has a bias towards x , whereas for larger n , it is close to $\hat{\mu}_n$. A similar effect can be observed for v_0^* : the impact of x vanishes when n grows, and for larger n it will be approximately equal to \hat{v}_n . Along the same lines it can be argued that for large n the approximation (17) decreases to

$$\exp\left(-\frac{1}{2} \frac{(x - \hat{\mu}_n)^2}{\hat{v}_n}\right).$$

In other words, when n grows large, the uncertainty of the estimates $\hat{\mu}_n$ and \hat{v}_n decreases, and, as a result, the approximation $\Delta_{\text{IS}}(x)$ converges to the certainty-equivalence-based estimate (15). So, in general the frequentist and Bayesian approaches lead to different approximations, where the Bayesian approach is the more conservative.

Remark 5.1 In the context of loss-probability-based utility functions, notice that the performance measure of loss discussed in this section, $\Delta(x)$, is based on the *total* allocated bandwidth. This is in contrast to the loss probability $\delta(x)$ based on *excess* bandwidth (see Equation (11)). To use our estimators in the iterative scheme for allocating excess bandwidth, we can use the following:

$$\delta_{\text{CE}}(x) = \Delta_{\text{CE}}(x + \hat{\mu}_n), \quad \text{and} \quad \delta_{\text{IS}}(x) = \Delta_{\text{IS}}(x + \hat{\mu}_n).$$

In doing this, the Inverse-Sanov (IS) estimator has an element of certainty-equivalence to it, in that we use the sample mean $\hat{\mu}_n$ as the estimate of the true mean. However, as we will see through our numerical example, using this Inverse-Sanov estimator still yields a more conservative outcome.

Remark 5.2 The Inverse-Sanov methodology (in conjunction with the principle of the largest term) is a powerful tool in many other situations, and nicely captures the possible ‘dangers’ of using the CE-approach. Although it is not entirely in the scope of the present paper, we have decided to include here an example that relates to measurement-based admission control, as it nicely illustrates the advantages of the IS-approach.

Consider a link in a network, on which i.i.d. sources are active. Let each individual source transmit traffic at a rate that is distributed as a random variable X , which we assume to represent a Normal random variable, but with *unknown* mean and variance. Suppose we are able to obtain n measurements of this traffic rate (which we assume to be independent, and distributed as the random variable X), say, a_1, \dots, a_n . The link has service capacity $C \equiv nc$, and we focus on the regime of large n (i.e., we have many observations at our disposal).

We first compute the number of sources that can be admitted for the CE-approach, under the requirement that the probability of the aggregate input rate exceeding C is below ϵ , based on a_1, \dots, a_n . Denoting by the X_i the rates of the individual sources, we wish to determine the biggest m such that

$$\mathbb{P} \left(\sum_{i=1}^m X_i \geq nc \right) \leq \epsilon.$$

As we consider certainty equivalence, we assume that the X_i follow a Normal distribution with mean $\hat{\mu}_n$ and variance \hat{v}_n . Applying the usual large-deviations approximation, we see that to find m , we must solve

$$\exp \left(-\frac{1}{2} \frac{(nc - m\hat{\mu}_n)^2}{m\hat{v}_n} \right) = \epsilon.$$

Elementary calculus yields that

$$m \equiv m(n) = \left(\frac{c}{\hat{\mu}_n} \right) n - \left(\frac{\sqrt{\hat{v}_n}}{\hat{\mu}_n} \sqrt{\frac{c}{\hat{\mu}_n}} \cdot \sqrt{-\log \epsilon} \right) \sqrt{n} + o(\sqrt{n}). \quad (18)$$

We now approximate the probability that the aggregate input rate exceeds C is below ϵ under the CE-based admission control. Suppose that we accept these m sources. It is natural to ask whether the probability of the aggregate input rate exceeding C is indeed in the order of ϵ ; if it is, then the CE-approach is a reliable method. Given our discussion on CE, one might expect that the probability is larger than ϵ . To verify this, we follow the Bayesian approach:

$$\mathbb{P} \left(\sum_{i=1}^m X_i \geq nc \right) \approx \int_{\mu_0, v_0} \exp \left(-\frac{1}{2} \frac{(nc - m\mu_0)^2}{mv_0} - \frac{n}{2} \left(\log \left(\frac{v_0}{\hat{v}_n} \right) + \frac{(\hat{\mu}_n - \mu_0)^2 + \hat{v}_n - v_0}{v_0} \right) \right) d\mu_0 dv_0.$$

Again applying the principle of the largest term, we find that the optimizing μ_0 and v_0 are

$$\mu_0^* = \frac{n}{m+n}(c + \hat{\mu}_n); \quad v_0^* = \frac{1}{m} \cdot \frac{1}{n+m}(nc - m\hat{\mu}_n)^2 + \hat{v}_n,$$

which yields the approximation, with $m \equiv m(n)$ as given in (18),

$$\mathbb{P}\left(\sum_{i=1}^m X_i \geq nc\right) \approx \left(1 + \frac{1}{m} \cdot \frac{1}{m+n} \cdot \frac{(nc - m\hat{\mu}_n)^2}{\hat{v}_n}\right)^{-n/2}.$$

According to (18) it holds that $(nc - m\hat{\mu}_n)^2 = \hat{v}_n \cdot (c/\hat{\mu}_n) \cdot \sqrt{-\log \epsilon} \cdot n + o(n)$, i.e., essentially linear in n . With $m \approx nc/\hat{\mu}_n$, we obtain that for n large

$$\begin{aligned} \mathbb{P}\left(\sum_{i=1}^m X_i \geq nc\right) &\approx \exp\left(-\frac{1}{2} \cdot \frac{\hat{\mu}_n}{\hat{\mu}_n + c} \cdot (-\log \epsilon)\right) = (\sqrt{\epsilon})^{\hat{\mu}_n/(\hat{\mu}_n + c)} \\ &= (\sqrt{\epsilon})^{\hat{\mu}_n/(\hat{\mu}_n + c/n)}. \end{aligned}$$

Since for large n we have that $\hat{\mu}_n/(\hat{\mu}_n + c/n) \approx 1$, we observe that this probability is approximately $\sqrt{\epsilon}$, which is considerably larger than ϵ .

We may draw the following conclusions from the above computations. In the first place, consider CE-based admission control, i.e., admit m connections, with m as in (18). The performance target is that the probability of exceeding C is below ϵ , but it turns out that this probability is rather in the order of $\sqrt{\epsilon}$, i.e., considerably too high. Apparently, the uncertainty of the estimates $\hat{\mu}_n$ and \hat{v}_n is so strong that the CE-based admission control is too optimistic.

In the second place, we observe that, in order to meet the performance target, one may still use the CE-method, but now with an adjusted ϵ , i.e., $\epsilon' := \epsilon^2$. These observations are in line with the findings of [7, Section 2]; notice that the approaches used are entirely different. \diamond

Example of our proposed loss estimators. Figure 1 shows the difference between the log of the two loss estimators. Using a sample size of $N = 50$, we consider a system with traffic source parameters $\mu = 2$, $\sigma = 0.15$ and $T = 1$. The graph is generated by taking 20 samples from a $\mathcal{N}(2, 0.15^2)$ distribution, constructing the estimators $\Delta_{\text{CE}}(x)$ and $\Delta_{\text{IS}}(x)$, respectively, and plotting the log of the means of the calculated quantities.

The graph shows the log of the two estimators of loss as the allocated capacity x increases. The upper curve is the result of the IS estimator, confirming our previous observation that it was the more cautious estimator of the loss probability.

6 Numerical example

We illustrate the performance of the two loss estimators via a numerical example. Consider a network of four links of capacity $C = 20$, shared by five ISPs, of the form given in




Figure 1: Comparison of CE and IS loss estimators with $N = 50$; System parameters: $\mu = 2$, $\sigma = 0.15$, $T = 1$.

Figure 2. The routes are denoted by dashed lines, and traverse at least one link between an origin and a destination node.




Figure 2: Network of four links shared by five routes

One ISP traverses 3 links, two ISP's traverse 2 links and the remaining ISP's traverse a single link. We assume each ISP has a utility function of the form

$$\bar{U}(\delta) = \alpha - \frac{\beta}{(1 - \delta)}. \quad (19)$$

As seen in Example 4.1, this utility function induces a utility function in terms of excess bandwidth that is strictly increasing and concave. Given the form of (19), the demand for excess bandwidth tends to 0 as the price $p \rightarrow \infty$. This implies that in this example we know *a priori* that $p^* = p_{\max}$. The numerical example uses network parameters given in Table 1.

Network Parameter	Parameter Value
Links	$c = [20, 20, 20, 20]$
Utility function	$\alpha = 1, \beta = 10$
True mean	$\mu = [3, 3, 3, 3, 3]$
True standard deviation	$\sigma = [0.2, 0.2, 0.2, 0.2, 0.2]$
Price cap	$p_{\max} = 1$ for all links
Step-size	$\epsilon = 0.02$

Table 1: Network parameters for numerical example.

Updates to the link prices and multipliers are made according to equations (9)-(10). From these equations, it can be seen that, at each link, not only the demand for excess bandwidth and the corresponding derivative should be known, but also the mean traffic rate. In the measurement-based context, the ISP s must estimate these quantities based on observations on the system. At each iteration, an independent sample of size n is taken. From this sample, the empirical mean can be calculated, and the ISP s are first allocated this quantity as an estimate of the mean rate (the minimum amount required to carry their traffic). This is aggregated on each link ℓ to estimate $\bar{\mu}_\ell$, the total mean rate on link ℓ , leaving a remaining $c_\ell - \hat{\mu}_\ell$ units of bandwidth available.

The amount of this remaining bandwidth that gets allocated to the ISP s is derived from the demand for excess bandwidth. This depends on how δ is evaluated, which brings us to our two loss estimators of Section 5. As explained, the Inverse-Sanov estimator $\Delta_{\text{IS}}(x)$ is more conservative than the Certainty-Equivalence estimator $\Delta_{\text{CE}}(x)$, in the sense that for a given bandwidth x , the estimated loss probability is higher using $\Delta_{\text{IS}}(x)$.

Since we know *a priori* that in this example $p^* = p_{\max}$ for all links, it follows that the excess bandwidth demanded by the ISP s is larger when they use $\Delta_{\text{IS}}(x)$ as the estimator of the loss probability. The final profit $\pi(p^*)$ is made up of the link prices multiplied by the total allocated bandwidth, where the latter is made up of the mean rate plus the demand for excess bandwidth. Hence, the previous discussion implies that the profit will also be larger when $\Delta_{\text{IS}}(x)$ is used.

Figure 3 shows the profit derived as the iterative procedure progresses, when a sample size of $n = 20$ is used. Profit is calculated as the current price multiplied by the estimated mean rate plus the demand for excess bandwidth at the prevailing price. The solid straight lines indicate the average behavior of the IS and CE estimators, respectively. As expected, on average, the profit derived using $\Delta_{\text{IS}}(x)$ is larger than when using $\Delta_{\text{CE}}(x)$. The difference between the profits can be thought of as a premium that the network earns from being used by risk-averse ISP s.




Figure 3: Derived profit using two estimators of loss and the first-order iterative method.

The difference between the IS and CE estimators is further highlighted by Figure 4. It shows the total allocated capacity for the three-link route as the price update procedure progresses. To aid comparison between the scenarios, we have removed the fluctuations in bandwidth as the iterative scheme progresses and have only shown the ‘average’ behavior (achieved by assuming the ISP views a sample mean $\hat{\mu}_n$ and sample variance \hat{v}_n equal to the true mean and variance at each iteration). The benefit of this is the ease with which we can see how sample size affects the impact of the Inverse-Sanov estimator. The bottom-most curve is the total allocated bandwidth when the ISP uses the CE estimator. When using measurements, the $\Delta_{\text{CE}}(x)$ -based total allocation fluctuates about this curve. So, quite frequently, the CE-based allocation is below the true optimal allocation. The remaining curves are the total allocated bandwidth when using the IS estimator for various sample sizes. The uppermost is when the smallest sample size $n = 10$ is used. Here the ISP mistrusts the observed mean and variance, and hence compensates by demanding more bandwidth. This mistrust diminishes as n is made larger (here $n = 20, 50, 100$ are shown), which translates to a corresponding decrease in the profit to the network.

7 Concluding remarks and future work

This paper studied the workings of a network in which the network and the users of the network have misaligned objectives. In particular, we considered the situation where each link is able to set a price for bandwidth and the users vary their demand depending

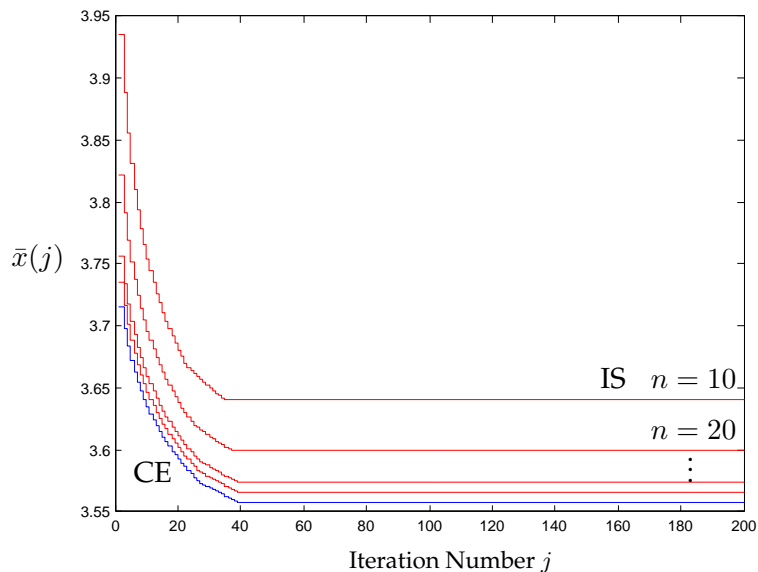


Figure 4: Total allocated capacity to ISP 1 in the network – includes mean rate + demand for excess bandwidth.

on the prevailing price. Their demand is derived from a utility function, used to measure how greatly they value bandwidth. In our context, the users can be thought of as ISPs competing for the finite bandwidth on each link of the network. In economic terms, the links are price setters and have the objective of maximizing their profit, while the users are price takers and have the objective of maximizing their compensated utility (utility less the cost of acquiring bandwidth). This can be contrasted with the network situation in which the network and the users have aligned objectives, known as social welfare maximization.

The first contribution of this paper was a provably convergent iterative update scheme for a single link model. The procedure updated the price for bandwidth so the network could achieve its maximum profit, while satisfying the link capacity constraints. It relied on communication between each user and the link, but not between the users. It thus satisfied the requirement of being distributed; this is often essential in the network context. Users supplied the link their demand for bandwidth, as well as a measure of the elasticity of their demand, at each iteration. This information was used by the link to optimally increment the price.

The second contribution of the paper was to provide a way of using the derived iterative procedure in a measurement-based context. This is required since each ISP is used by a large number of end-users. The desire for bandwidth of these end-users fluctuates over time and hence an ISP does not know its utility function with certainty. Instead it must

base its decisions on observation of the system. We proceeded using the assumption of Gaussian input traffic, appropriate here because of the high aggregation of end-users. In this context, we derived an estimator of loss probability that is useful for risk-averse ISPs. It was shown how this estimator could be used in the iterative procedure for bandwidth and price updates, and we illustrated the modified update algorithm via a simple numerical example.

The work contained here can be extended in several ways. First, the set up of the profit maximization problems could be made more complex. This may be desirable if, for example, users' demand for bandwidth depended not only on the prevailing price, but also on some measure of congestion present on the link. We anticipate that this would make the iterative scheme more complicated. This is especially so given the requirement that the scheme be distributed, and hence the ISPs cannot directly communicate with each other. Another interesting possibility is to investigate the scenario where there exist multiple network owners; hence the market for bandwidth is an oligopoly, rather than a monopoly.

Instead of making the profit maximization more complex, it would also be of interest to further investigate the convergence properties of the iterative scheme in the measurement-based context. This may require use of stochastic approximation techniques and a more in-depth analysis of the noise associated with the estimated quantities.

Appendix

This appendix contains details of the convergence of a first-order update scheme applied to the augmented Lagrangian of an inequality constrained optimization problem, see Algorithm 3.1. A key element in the proof are the update formulae (20)-(21); it is a straightforward exercise to verify that these are in line with the update formulae of Algorithm 3.1, i.e., (7)-(8)

Our proof follows a similar format to Bertsekas [1, Prop. 4.4.2], which proves convergence of a first-order scheme applied to the conventional Lagrangian of an equality constrained optimization problem.

Consider the optimization problem:

$$\max_x J(x) \quad \text{subject to} \quad g(x) \leq 0,$$

for which the SOSOC holds true, and the optima x^* are regular. Let η denote the vector of multipliers of dimension L and let R denote the dimension of the x . Consider the augmented Lagrangian with penalty parameter m (cf. $\mathcal{L}_m(p, \eta)$ introduced in Section 3), so that the function is locally concave at (x^*, η^*) .

Let x_j and η_j be the j -th iteration of the control and the multiplier. The first-order update algorithm can be described by $x_{j+1} = G(x_j, \eta_j)$ and $\eta_{j+1} = H(x_j, \eta_j)$, where

$$H(x, \eta) := \max\{0, \eta + \epsilon g(x)\}, \tag{20}$$

$$G(x, \eta) := x + \epsilon(\nabla J(x) - \nabla g(x) \max\{0, \eta + mg(x)\}); \tag{21}$$

here $\nabla J(x)$ and $\nabla g(x)$ denote derivatives with respect to x , and the max is taken element-wise. Equalities to follow are also taken element-wise.

According to [1, Prop. 4.4.1], the iterations converge to a fixed point, provided that the system starts in a neighborhood S such that the mappings G and H are continuously differentiable and the eigenvalues of Γ^* , a matrix of derivatives of G and H evaluated at the fixed point, lie within the unit circle. If the mappings G and H are continuously differentiable in S , then the entries of the matrix Γ are continuous. Since the eigenvalues of a square matrix depend continuously on the entries [1, Prop. A.14], for small ϵ and $(x, \eta) \in S$, the eigenvalues will remain within the unit circle.

It clear that (20) and (21) are both continuous functions (in x and η), but are not continuously differentiable. $G(x, \eta)$ has a discontinuity in its derivative at points where $\eta = -mg(x)$, while $H(x, \eta)$ has a discontinuity in its derivative where $\eta = -\epsilon g(x)$. These discontinuities mean that we cannot immediately apply Proposition 4.4.1 [1], as is possible in the equality-constrained case, to prove convergence.

We observe that away from these discontinuities, the mappings *are* continuously differentiable. So, if we can establish that the optimal solution (x^*, η^*) is not at a discontinuity, then there exists an open set around (x^*, η^*) where G and H are continuously differentiable.

At the optimal solution, there is a set of constraints that are active ($g_j(x^*) = 0$, for j in the active set $A(x^*)$) and a set of constraints that are inactive ($g_j(x^*) < 0$ for $j \notin A(x^*)$). By strict complementarity, for $j \in A(x^*)$ we have $\eta_j > 0$ and for $j \notin A(x^*)$ we have $\eta_j = 0$. Therefore, $\eta^* \neq -mg(x^*)$ and $\eta^* \neq -\epsilon g(x^*)$. Hence, the solution (fixed point of the updates) is not at a discontinuity and therefore there is a neighborhood around (x^*, η^*) for which the mappings are continuously differentiable. We thus consider Γ^* in this neighborhood and show that its eigenvalues lie within the unit circle.

Let I_n denote the identity matrix of dimension $n \in \mathbb{N}$. For (x, η) in the neighborhood around (x^*, η^*) for which the mappings are continuously differentiable, with a slight abuse of notation with respect to vectors,

$$\begin{aligned}\nabla_x G(x, \eta) &= I_R + \epsilon \nabla_{xx}^2 \mathcal{L}_m(x, \eta) \\ \nabla_\eta G(x, \eta) &= \begin{cases} -\epsilon \nabla g_j(x) & \text{if } j \in A(x^*) \\ 0 & \text{otherwise} \end{cases} \\ \nabla_x H(x, \eta) &= \begin{cases} \epsilon \nabla g_j(x) & \text{if } j \in A(x^*) \\ 0 & \text{otherwise} \end{cases} \\ \nabla_\eta H(x, \eta) &= \begin{cases} 1 & \text{if } j \in A(x^*) \\ 0 & \text{otherwise.} \end{cases}\end{aligned}$$

Now consider the matrix of derivatives Γ^* . Let the number of active constraints be denoted k , and $\nabla_{Ag}(x)$ denote the $R \times k$ matrix containing the constraint gradients (with respect to x) for these active constraints. Then Γ^* is given by

$$\Gamma^* = \begin{bmatrix} I_R + \epsilon \nabla_{xx}^2 \mathcal{L}_m(x^*, \eta^*) & \epsilon \nabla_{Ag}(x^*) & \mathbf{0} \\ -\epsilon \nabla_{Ag}(x^*)^T & I_k & \mathbf{0} \\ \mathbf{0} & \mathbf{0} & \mathbf{0} \end{bmatrix}, \quad (22)$$

where $\mathbf{0}$ denotes a matrix of the appropriate dimension. The difference between Γ^* for the inequality and the equality cases is that for the latter, there are no inactive constraints. Hence the

block structure of Γ^* is more complicated here. The matrix Γ^* given in (22) can be written as $\Gamma^* = I_{R+L} - \epsilon B$, where

$$B := \begin{bmatrix} -\nabla_{xx}^2 \mathcal{L}_m(x^*, \eta^*) & -\nabla_{Ag}(x^*) & \mathbf{0} \\ \nabla_{Ag}(x^*)^T & \mathbf{0} & \mathbf{0} \\ \mathbf{0} & \mathbf{0} & \epsilon^{-1} I_{L-k} \end{bmatrix}.$$

If we can show the real part of each eigenvalue of B is strictly positive, then we can use ϵ as a ‘scaling’ factor to ensure the eigenvalues lie within the unit circle. We are then done.

Let the vector \hat{y} denote the complex conjugate of the vector y . Let β be an eigenvalue of B and let $(z, w, v) \neq 0$ be the corresponding eigenvector, where z , w , and v are complex vectors of dimension R , k , and $L - k$, respectively. We have

$$\operatorname{Re} \left\{ (\hat{z}^T, \hat{w}^T, \hat{v}^T) B \begin{pmatrix} z \\ w \\ v \end{pmatrix} \right\} = \operatorname{Re} \left\{ \beta (\hat{z}^T, \hat{w}^T, \hat{v}^T) \begin{pmatrix} z \\ w \\ v \end{pmatrix} \right\} = \operatorname{Re}(\beta) (\|z\|^2 + \|w\|^2 + \|v\|^2), \quad (23)$$

by the definition of β being an eigenvalue of B , while using the form of B , we have

$$\begin{aligned} & \operatorname{Re} \left\{ (\hat{z}^T, \hat{w}^T, \hat{v}^T) B \begin{pmatrix} z \\ w \\ v \end{pmatrix} \right\} \\ &= \operatorname{Re} \left\{ -\hat{z}^T \nabla_{xx}^2 \mathcal{L}_m(x^*, \eta^*) z + \hat{w}^T \nabla_{Ag}(x^*)^T z - \hat{z}^T \nabla_{Ag}(x^*) w + \hat{v}^T \epsilon^{-1} I_{L-k} v \right\} \\ &= \operatorname{Re} \left\{ -\hat{z}^T \nabla_{xx}^2 \mathcal{L}_m(x^*, \eta^*) z + \hat{v}^T \epsilon^{-1} I_{L-k} v \right\}, \end{aligned} \quad (24)$$

which follows, since for any real $R \times L$ matrix D , we have $\operatorname{Re}\{\hat{z}^T D^T w\} = \operatorname{Re}\{\hat{w}^T D z\}$.

Combining equations (23) and (24) yields

$$\operatorname{Re} \left\{ -\hat{z}^T \nabla_{xx}^2 \mathcal{L}_m(x^*, \eta^*) z + \hat{v}^T \epsilon^{-1} I_{L-k} v \right\} = \operatorname{Re}(\beta) (\|z\|^2 + \|w\|^2 + \|v\|^2). \quad (25)$$

It is left to prove that the left-hand side of (25) is positive for $(z, w, v) \neq 0$, since this implies that $\operatorname{Re}(\beta) > 0$, which is what we set out to show.

If $z \neq 0$ and/or $v \neq 0$, then the left-hand side of equation (25) is strictly positive. This is because $\mathcal{L}_m(x^*, \eta^*)$ is locally concave, hence $-\hat{z}^T \nabla_{xx}^2 \mathcal{L}_m(x^*, \eta^*) z > 0$ for $z \neq 0$. In addition, $\epsilon^{-1} \hat{v}^T I_{L-k} v > 0$ for $v \neq 0$.

If $z = 0$ and $v = 0$, by the assumption that the eigenvector is not a vector of zeroes, this implies that $w \neq 0$. Then, the left-hand side of (25) is zero and hence $\operatorname{Re}(\beta) = 0$. Therefore, the equation

$$B \begin{pmatrix} z \\ w \\ v \end{pmatrix} = \beta \begin{pmatrix} z \\ w \\ v \end{pmatrix}$$

reduces to $\nabla_{Ag}(x^*) w = 0$. Since x^* is regular, the active constraint gradients are linearly independent. Hence, $w = 0$. This yields a contradiction. Therefore, (25) is strictly positive, which implies $\operatorname{Re}(\beta) > 0$. This is what we set out to show. Therefore, for $\epsilon > 0$, there exists a neighborhood S of (x^*, η^*) such that if $(x_0, \eta_0) \in S$, then the first-order updates are a contraction mapping and converge to (x^*, η^*) as $j \rightarrow \infty$.

References

- [1] D. P. Bertsekas, *Nonlinear Programming*. Second Edition, Athena Scientific, Belmont MA, USA, 1999.
- [2] N. G. Duffield, "A Large Deviation analysis of errors in measurement based admission control to buffered and bufferless resources", *Queueing Systems*, **34**: 131–168, 2000.
- [3] C. Fraleigh, F. Tobagi, and C. Diot, "Provisioning IP Backbone Networks to Support Latency Sensitive Traffic", *In Proceedings IEEE INFOCOM 2003*, San Francisco CA, USA, 2003.
- [4] E. W. Fulp, and D. S. Reeves, "Optimal Provisioning and Pricing of Internet Differentiated Services in Hierarchical Markets", *In ICN 2001, LNCS 2093*, P. Lorenz (Ed.) 409–418, 2001.
- [5] A. Ganesh, N. O'Connell and D. Wischik, *Big Queues, Lecture Notes in Mathematics 1838*. Springer, Berlin, 2004.
- [6] A. Ganesh and N. O'Connell, "An inverse of Sanov's theorem", *Statistics and Probability Letters*, **42**: 201–206, 1999.
- [7] M. Grossglauser and D. Tse, "A Framework for Robust Measurement-Based Admission Control", *IEEE/ACM Transactions on Networking*, **7**: 293–309, 1999.
- [8] X. Huang, A. Ozdaglar, and D. Acemoğlu, "Efficiency and Braess' Paradox under Pricing in General Networks", *IEEE Journal on Selected Areas in Communications*, **24**: 977–991, 2006.
- [9] R. Johari and D. Tan, "End-to-end congestion control for the Internet: delays and stability", *IEEE/ACM Transactions on Networking*, **9**: 818–832, 2001.
- [10] F. P. Kelly, "Charging and rate control for elastic traffic", *European Transactions on Telecommunications*, **8**: 33–37, 1997.
- [11] F. P. Kelly, A. K. Maulloo, and D. K. Tan, "Rate control for communication networks: shadow prices, proportional fairness, and stability", *Journal of the Operational Research Society*, **49**: 237–252, 1998.
- [12] J. Kilpi and I. Norros, "Testing the Gaussian approximation of aggregate traffic", *In Proceedings of Internet Measurement Workshop*, Marseille, France, 2002.
- [13] S. H. Low and D. E. Lapsley, "Optimization Flow Control, I: Basic Algorithm and Convergence", *IEEE/ACM Transactions on Networking*, **7**: 861–875, 1999.
- [14] D. G. Luenberger, *Linear and Nonlinear Programming*. Second Edition, Addison Wesley, Reading Mass., 1984.
- [15] R. van de Meent, M. Mandjes, and A. Pras, "Gaussian traffic everywhere?", *Proceedings IEEE International Conference on Communications*, Istanbul, Turkey, 2006.
- [16] J. Mo, and J. Walrand, "Fair End-to-End Window-Based Congestion Control", *IEEE ACM Transactions on Networking*, **8**: 556–567, 2000.
- [17] A. E. Ozdaglar, "Price Competition with Elastic Traffic", *Submitted*, 2006.
- [18] A. C. Pigou, *The Economics of Welfare*. Third Edition, Macmillan, London, 1929.
- [19] T. Roughgarden, *Selfish Routing and the Price of Anarchy*. MIT Press, Cambridge MA, USA, 2004.
- [20] H. van den Berg, M. Mandjes, R. van de Meent, Aiko Pras, F. Roijers, and P. Venemans, "QoS-aware bandwidth provisioning for IP network links", *Computer Networks*, **50**: 631–647, 2006.
- [21] G. Vinnicombe, "On the stability of networks operating TCP-like congestion control", *In Proceedings of the 15th IFAC World Congress on Automatic Control*, Barcelona, Spain, 2002.

Author Biographies

Michel Mandjes

Michel Mandjes received his M.Sc. (in both mathematics and econometrics) and Ph.D. degrees from the Vrije Universiteit (VU), Amsterdam, the Netherlands. After having worked as a member of technical staff at KPN Research (Leidschendam, the Netherlands) and Bell Laboratories/Lucent Technologies (Murray Hill NJ, USA), as a part-time full professor at the University of Twente, and as department head at the Centre for Mathematics and Computer Science (CWI), Amsterdam, he currently holds a full professorship at the University of Amsterdam, the Netherlands. His research interests include performance analysis of communication networks, queueing theory, Gaussian traffic models, traffic management and control, and pricing in multi-service networks.

Maya Ramakrishnan

Maya Ramakrishnan received her B.Sc(Hons)/B.Com (in both mathematics and finance) and Ph.D. degrees from The University of Melbourne, Australia. She has held a teaching position within the Department of Mathematics and Statistics at The University of Melbourne. She currently holds a position of post-doctoral research fellow within the Australian Research Council Centre of Excellence for Mathematics and Statistics of Complex Systems at The University of Melbourne. Her research interests include distributed control and optimization-based pricing in communication networks, and measurement-based schemes.