

# A Packet Marking Algorithm for Congestion Pricing

Aleksi Penttinen and Jorma Virtamo

Networking Laboratory

Helsinki University of Technology, P.O.Box 3000 FIN-02015 HUT, Finland

Email: {Aleksi.Penttinen, Jorma.Virtamo}@hut.fi

**Keywords:** Packet marking, congestion pricing, congestion management, model based control, MDP theory.

## Abstract

In the resource pricing concept, developed by Kelly et al. [1], the network marks packets with appropriate price signals to direct the users' actions. The theoretically correct congestion dependent prices, however, are not all known at the time when the marking occurs at a network resource. In this paper, we propose a packet marking algorithm which maximally utilises the information available at the time of marking. The known prices are marked directly on the packets, while the unknown ones are replaced by their expected values. Calculation of the expected values depends on a suitable system model. We derive these prices for a simple M/M/1/K resource model and generalise the result for GI/GI/1/K models using an approximate diffusion approach.

## INTRODUCTION

Intense research is going on to find solutions enabling concurrent transport of audio, video, and other real-time, as well as non-real-time data in the Internet, in a manner that satisfies various customer requirements. One of the most elegant proposals so far is the *resource pricing* [2, 3, 1, 4, 5] approach, suggesting the use of a simple feedback loop to provide stability, fairness and differentiated services for the network.

The idea of resource pricing is to put a price on the social cost of congestion occurring at the network resources and distribute it fairly to the users responsible for the detriments. The price acts as a sole incentive for the users to co-operate and avoid congestion, while the core of the network is left with the duty to signal the congestion costs back to users. In practice, the users are maximising the net utility of the whole network in distributed manner using TCP like flow control algorithms; the end nodes may behave as they wish, knowing that they will be charged accordingly. The signalling can be handled using a method similar to the IETF's ECN proposal [6] piggy-backing the congestion information on the packets using a single bit in the IP-packet header, i.e. *marking the packet*.

In [1] Gibbens and Kelly presented a connection between the packet level marks and the congestion costs as follows. Consider a queuing model for a buffer of an Internet resource

and assume overflowing packets to be the congestion cost to the system. Then, the users get fair prices if the network marks all the packets arriving during the critical congestion interval, i.e. the time between the start of the busy period and the last packet loss. The *sample path shadow price (SPSP)* of a packet is defined to equal one, if deleting the packet causes one less packet drop at the resource. Although the definition is simple, it poses a problem: some of the packets may have escaped from the system during a critical congestion interval before any packet losses have occurred. Therefore, alternative ways of determining the prices have to be considered.

Gibbens and Kelly suggested also two heuristic marking methods. First, one could keep count on the arriving packets from the start of the busy period and after an overflow mark a correct number of the leaving packets. Alternatively the marking could continue always after an overflow until the resource becomes empty. They saw it desirable, however, that the marking would take place somewhat earlier during the busy period so that the overflow itself could possibly be avoided.

Their second approach, called the virtual queue [1, 4], was developed for this early warning feature. In this scheme the resource maintains a virtual queue which has exactly the same arrival process as the real buffer, but the service rate is scaled down. This way the fictitious queue can anticipate the overflow behaviour of the real one. Marking is performed in this virtual queue using, e.g. a marking threshold.

Other marking methods, originally emerged in the ECN context (e.g. RED [7] and REM [8]), can also be used for resource pricing. These methods, however, base the marking decisions on the average queue size, which is generally too slow in reacting to the random fluctuations in the traffic.

A common issue with all these proposals is fairness, as discussed in [9]. In other words, the schemes are not concerned about generating exactly correct prices for the end-users as long as the method works on the flow level. Fairness is sacrificed for stability, as is the case often in engineering.

In this presentation we propose a packet marking algorithm which maximally utilises the available information at the time of marking. In this way both objectives, correct prices and an early warning of congestion, are achieved. In addition to the deterministic prices (for any packet for which such is available), we use the expected value of the price, the current system price, which is calculated for the unmarked packets upon

departure from the resource. Although the approach is model dependent, it is important to note that only the independence of inter-arrival and service time distributions is required for a general approximate solution.

The organization of the rest of the paper is as follows. The second section reviews the optimization model for resource pricing by Kelly et al. to shed light on the foundations of marking. In the third section we present a marking algorithm which is based on the expected prices. In particular, we solve these prices analytically for the M/M/1/K model and for the approximate diffusion model. Fourth section discusses a few numerical experiments before some summarising remarks are presented in the last section.

## MODEL

Here we briefly present the optimization model developed by Kelly et al. in a series of papers [2, 3, 1]. For more general overviews of the whole body of work the interested reader is referred to e.g. [10, 11, 12, 13].

Assume a set of resources ( $J$ ) and routes ( $R$ ) indexed by  $j$  and  $r$ , respectively. Each route, a subset of the resources  $r \subset J$ , can be seen as a user who has a transmission rate  $x_r$ . The flow through the resource  $j$  can be written as

$$y_j = \sum_{r \ni j} x_r. \quad (1)$$

Each user  $r$  has a utility denoted by  $U_r$  depending on the rate  $x_r$ . Suppose that the functions  $U_r(x_r)$  are concave, continuously differentiable, with  $U'_r(x_r) \rightarrow \infty$  as  $x_r \downarrow 0$  and  $U'_r(x_r) \rightarrow 0$  as  $x_r \uparrow \infty$ .

Assume now that a cost at rate  $C_j(y_j)$  incurs (to the system) at the resource  $j$  when the flow through it is  $y_j$ . The cost is here taken to be the number of lost packets, but it could alternatively depend on, e.g. the delay. Assume also that the users are maximising their own utilities (which are assumed to be additive to the cost). Then the system faces the problem of maximising the total net utility

$$\max_{\mathbf{x}} \sum_{r \in R} U_r(x_r) - \sum_{j \in J} C_j(y_j), \quad (2)$$

where the cost serves actually as a penalty function. To relate it to the congestion signals write

$$C_j(y_j) = \int_0^{y_j} p_j(\xi) d\xi, \quad (3)$$

in which the value of the function

$$p_j(y_j) = \frac{d}{dy_j} C_j(y_j), \quad (4)$$

is called the *shadow price* of the resource  $j$ . It can be seen as the rate of feedback signals generated by the resource  $j$ . If the users are maximising their utilities and the feedback signals

are “charged” from the users, the system evolves towards the social optimum [3], i.e. to the unique solution of the problem (2).

One can easily see that from the stability point of view the requirements on the pricing method are rather relaxed. The solution of the optimization problem (2) can be achieved with only weak conditions on the penalty functions  $C_j(y_j)$ .

However, maximum efficiency is reached only by correct price information and a true service differentiation is achieved only by means of fair pricing. Within a mathematical model these objectives can be achieved conveniently by a pricing method which provides an early warning of congestion at the same time.

## SYSTEM PRICES

The ideal SPSP scheme could be replaced by marking packets with the probability of the packet belonging to a critical congestion interval. This is equal to the probability that a later packet will overflow during the same busy period, as was noted already in [1]. This expected value of the price is well defined for the system at any time and we refer to it as the *system price*. Our packet marking scheme consists of two rules:


- mark all the packets in the buffer with the price 1 when an overflow occurs,
- mark all the non-marked packets as they are leaving the buffer with the (real valued) price of the system immediately after the departure (to get the latest information on the system price).

When using the one bit ECN style marking, the system price corresponds the marking probability. The users have to estimate the prices from the flow of marks, which slows down the response to altering prices. If all the users have the same packet length distribution, the system price can be determined straightforwardly for various queuing models as the following example shows.

### Simple Queuing Model

Consider an M/M/1/K queue model for a network resource. Packets arrive independently according to a Poisson process with the parameter  $\lambda$  and the service times are exponentially distributed with the parameter  $\mu$ . Denote the traffic intensity  $\rho = \lambda/\mu$ . From the user point of view we shall make no difference between whether the packet was discarded or merely just given the price 1.

First, for purposes of determining the probability of overflow from a given state before the end of the busy period an embedded Markov chain, so called jump chain, is constructed. The states of this chain are the queue occupancies after any change (arrival or departure) in the queue. The chain has the



**Figure 1.** Embedded Markov chain of the M/M/1/K-queue.

transition probabilities

$$u = \frac{\lambda}{\lambda + \mu} = \frac{\rho}{\rho + 1}, \quad (5)$$

$$d = \frac{\mu}{\lambda + \mu} = \frac{1}{\rho + 1}, \quad (6)$$

upwards and downwards, respectively, and absorbing states at 0 and  $K + 1$ , see Figure 1.

The interesting event here is the absorption to state  $K + 1$  (overflow), as opposed to state 0, when starting from a given state of the system. To this end, denote the probability of absorption to state  $K + 1$  from state  $n$  by  $p_n$ . Due to the Markovian property of the chain, it holds that

$$\begin{aligned} p_n &= u p_{n+1} + d p_{n-1}, \quad n \in \{1, \dots, K\}, \\ p_0 &= 0, \quad p_{K+1} = 1. \end{aligned} \quad (7)$$

This linear homogenous difference equation can be solved by using standard methods [14] leading to the simple solution

$$p_n = \rho^{K+1-n} \frac{(\rho^n - 1)}{(\rho^{K+1} - 1)}. \quad (8)$$

This pricing formula can be shown to give the correct prices. Consider an M/M/1/K system described above with a number of users, indexed by  $r$ , with a Poisson arrival process with parameters  $\lambda_r$ . The aggregate arrival process to the resource is then also Poisson with mean  $\lambda = \sum \lambda_r$ . The service rates are exponentially distributed with the mean service time  $1/\mu$ . Denote the traffic intensity of user  $r$  by  $\rho_r = \lambda_r/\mu$  and the total intensity  $\rho = \lambda/\mu$ . The steady state probability of state  $i$  is

$$\pi_i = \rho^i \frac{1 - \rho}{1 - \rho^{K+1}}. \quad (9)$$

The expected rate of loss of data, the congestion cost, is given by arrivals to the system in state  $K$ , i.e. at the bit rate

$$C(\rho) = \rho P(\text{system full}) = \rho \rho^K \frac{1 - \rho}{1 - \rho^{K+1}}, \quad (10)$$

and thus the shadow price is obtained by differentiation

$$p(\rho) = C'(\rho) = \frac{\rho^K (\rho^{K+2} - (K+2)\rho + K+1)}{(\rho^{K+1} - 1)^2}. \quad (11)$$

Assume that upon arrival to the system, a packet is given the price determined by the equation (8) when *after* the arrival

the system is in state  $n$ . Assume also that the system is at equilibrium and so the expected rate of incoming price to the user  $r$  is

$$\begin{aligned} &\rho_r \sum_{i=0}^K \pi_i p_{i+1} \\ &= \rho_r \rho^K \frac{\rho - 1}{(\rho^{K+1} - 1)^2} \sum_{i=0}^K (\rho^{i+1} - 1) \\ &= \rho_r p(\rho). \end{aligned} \quad (12)$$

This means that by marking each arriving packet with the price (8) gives the desired information within this simple Markovian model. In similar fashion it is possible to calculate the prices for various queuing models having independently distributed inter-arrival and service times, which is shown in [15].

We can also interpret the price as the additional cost to the system due to the arriving packet. In the framework of Markovian Decision Processes (MDPs) this equals the relative cost of a state. In this spirit we derive the same result obtained above using a heuristic approach.

### Simple Model And State Dependent Costs


The heuristic derivation of the relative cost of states, for the most part, is similar to the approach introduced by Krishnan [16] for the M/M/K/K loss system in the context of dynamic state-dependent routing.

Consider a packet arriving into the M/M/1/K system with  $n$  packets present and the load  $\rho = \lambda/\mu$ . At that point there are two decision alternatives: either admit the packet into the system or reject it. The shadow price will arise from the comparison of these decision alternatives as shown below.

If the packet is accepted, the system will be at state  $n + 1$  after the arrival. Shadow prices are defined to be the cost increment for a marginal increase in load. Hence the interesting property is the *increase in cost*, i.e. the increase in the expected number of lost packets in the future, due to the admitted packet. What is the expected increase in cost when starting from the state  $n + 1$  instead of  $n$ ? (In MDP terms, we calculate the difference of the *relative costs* of the states  $n$  and  $n + 1$ , but it is not necessary to go further into the MDP terminology here).

The cost of a sample path representing the queue occupancy at a time  $t$  is the number of blocked packets before  $t$ . Consider two arbitrary paths,  $\omega_n$  starting from the state  $n$ , and  $\omega_{n+1}$  starting from state  $n + 1$ .

Due to the Markovian property of the process, from the instant onwards when  $\omega_n$  reaches the state  $n + 1$  for the first time (denote the time this happens, i.e. the first passage time, by  $t_n^*$ )  $\omega_n$  is statistically identical to  $\omega_{n+1}$ . Naturally, there cannot be any overflows on  $\omega_n$  before  $t_n^*$  as the process has to pass through  $n + 1$  in order to reach higher states and so also the blocking state. In other words, at any given time  $t > t_n^*$ ,  $\omega_n$  is stochastically identical to what  $\omega_{n+1}$  was the time  $t_n^*$  ago and so one can expect an equal number of overflows during  $(0, t)$



**Figure 2.** The effect of the starting state on overflows.

on  $\omega_n$  and  $(0, t - t_n^*)$  on  $\omega_{n+1}$ . Thus the increase in cost at the time  $t$ , the expected difference in numbers of overflows on the paths, is the expected number of lost packets of  $\omega_{n+1}$  between  $(t - t_n^*, t)$ . This is illustrated in Figure 2.

When  $t \rightarrow \infty$ , effects of the initial value vanish, the system is at equilibrium and the probability that the system is full (and all the arriving packets overflow) is given by

$$\text{Ovf}(K, \rho) = \frac{\rho^K}{\sum_{i=0}^K \rho^i} = \rho^K \frac{(1-\rho)}{(1-\rho^{K+1})}. \quad (13)$$

On average  $\lambda E[t_n^*]$  packets arrive during the time  $(t - t_n^*, t)$  and hence the increase in expected number of lost packets becomes

$$p_n = \lambda E[t_n^*] \text{Ovf}(K, \rho). \quad (14)$$

Finally,  $E[t_n^*]$  can be determined by considering an M/M/1/n system. Immediately after a packet overflows the system is full, i.e. in the state  $n$ . Next overflow in this system corresponds to the transition from state  $n$  into  $n+1$  in the M/M/1/K. Thus,  $t_n^*$  is distributed as the time between subsequent overflow events in an M/M/1/n system (arrivals into full system) and has the expectation equal to the inverse of the overflow rate

$$E[t_n^*] = \frac{1}{\lambda \text{Ovf}(n, \rho)}. \quad (15)$$

Now we can state our result: the increase in cost – the shadow price – is given by

$$p_n = \frac{\text{Ovf}(K, \rho)}{\text{Ovf}(n, \rho)}, \quad (16)$$

noting that  $n$  represents here the state *before* the packet is accepted into the system while the price (8) is calculated *after* the packet has been placed into the queue.

## Gaussian Model

By analysing the state-dependent overflow probabilities (or the relative costs) for other simple queuing models it can be seen that the *form* of the pricing function on the packet level remains roughly the same: under heavy traffic (offered load close to one) the pricing function is almost linear with respect to states, while light traffic causes prices to approach a step function. If we could capture this observation into a single

model and solve the prices, we would have a formula for pricing all the non-correlated traffic. For this purpose we examine a Gaussian traffic model, the diffusion approximation.

Diffusion approximation is based on the fact that many independent and identically distributed sources generate essentially a Gaussian type of aggregate behaviour due to the central limit theorem. The reason to use the continuous approximation is thus obvious; we have a theoretical support for the results regardless of the source distributions which otherwise would have to be modelled.

For any diffusion process there are two commonly defined infinitesimal parameters:

$$\mu(x, t) = \lim_{h \rightarrow 0} \frac{E[X(t+h) - X(t) | X(t) = x]}{h}, \quad (17)$$

$$\sigma^2(x, t) = \lim_{h \rightarrow 0} \frac{E[\{X(t+h) - X(t)\}^2 | X(t) = x]}{h}, \quad (18)$$

which are called the drift parameter and the infinitesimal variance, respectively. For our purposes it is sufficient to consider the time homogenous processes for which the infinitesimal parameters depend only on the state  $x$  and not on  $t$ .

Following the procedure in [17], we now set the state 0 and the predefined overflow limit  $c$  as absorbing barriers. The problem is to find the probability of absorption into  $c$  when starting from the position  $X(0) = x$ ,  $0 < x < c$ . Denote this probability by  $u(x)$  and select a small time duration  $h$ , so small that the absorption probability is negligible. Naturally one can set  $u(0) = 0$  and  $u(c) = 1$ . Further, let  $\delta X = X(h) - x$  and one can write

$$\begin{aligned} u(x) &= E[u(X(h)) | X(0)] + o(h) \\ &= E[u(x + \delta X) | X(0)] + o(h) \\ &= u(x) + E[\delta X | X(0)] u'(x) + \\ &\quad + \frac{1}{2} E[(\delta X)^2 | X(0)] u''(x) + o(h) \\ &= u(x) + \mu(x) h u'(x) + \\ &\quad + \frac{1}{2} \sigma^2(x) h u''(x) + o(h). \end{aligned} \quad (19)$$

From this, by subtracting  $u(x)$ , dividing by  $h$ , and taking the limit when  $h$  tends to zero, we get a differential equation for the desired probability.

$$\begin{aligned} 0 &= \mu(x) \frac{du}{dx} + \frac{1}{2} \sigma^2(x) \frac{d^2u}{dx^2}, \\ 0 &< x < c, \quad u(0) = 0, \quad u(c) = 1. \end{aligned} \quad (20)$$

Assume that  $\sigma^2(x) > 0$  and define the scale function of the process as

$$S(x) = \int^x \exp \left\{ - \int^\eta \frac{2\mu(\eta)}{\sigma^2(\eta)} d\eta \right\} d\tau. \quad (21)$$

Then the general solution of (20) can be written [17] as

$$u(x) = \frac{S(x) - S(0)}{S(c) - S(0)}, \quad 0 \leq x \leq 1. \quad (22)$$

The result given here can be easily applied to provide results for all time homogenous diffusion processes. For the simplest model, we follow the example of Harrison and Patel [18], and shed light on a GI/GI/1 queuing model with independent inter-arrival time  $A$  and independent service time  $B$  with means  $E[A] = 1/a$  and  $E[B] = 1/b$ , respectively. The corresponding variances are  $c_a^2/a^2$  and  $c_b^2/b^2$ , where the  $c_a$  and  $c_b$  are the corresponding variation coefficients.

Now we can approximate the queue state by a continuous path process  $X(t)$  with parameters which are actually constant in steady state. Using the notation above

$$\mu = \mu(x) = \lim_{t \rightarrow \infty} \mu(x, t) = a - b, \quad (23)$$

$$\sigma^2 = \sigma^2(x) = \lim_{t \rightarrow \infty} \sigma^2(x, t) = ac_a^2 + bc_b^2. \quad (24)$$

Thus, for a GI/GI/1 queue under the diffusion approximation we get

$$S(x) = C \exp\left(-\frac{2\mu x}{\sigma^2}\right) + D, \quad (25)$$

with constants  $C$  and  $D$ . Hence, the probability of overflow from state  $x$  becomes simply

$$u(x) = \frac{e^{-2\mu x/\sigma^2} - 1}{e^{-2\mu c/\sigma^2} - 1}, \quad (26)$$

where the barrier  $c$  is  $K + 1$  in the queuing model. Note that only the ratio of the drift and variance parameters has to be estimated.


Actually the result requires the mutual independence of the arrival and departure processes, which is not exactly true on large scale; there cannot be more departures than arrivals at any time and so the departure process is limited by the arrivals. However, during a busy period the inter-departure times are just the service times and hence independent from the arrival process.

The form of the price function can very well be approximated with the diffusion process absorption formula regardless of the model. However, the convergence of, e.g. processes with a deterministic component, towards a Gaussian process is rather slow at low traffic intensities. Therefore, we could include a single real valued correction parameter  $\delta \in (-1, 2)$  to the calculations reflecting the bias caused by the skewness of the counting process distribution. In other words, if the system is left to state  $n$  we evaluate the system price as  $u(n + \delta)$ .

More flexibility can be achieved by making the diffusion parameters state dependent in (22). For example, if there are users whose round trip times are comparable to the busy periods (i.e. they can react within a busy period), or if the number of active flows is small, arrivals in the higher states may be rarer than into a relatively empty system.

**Table 1.** Queuing models for comparison

Model	$\lambda$	$\mu$	$\delta$
M/M/1/25	3	4	0
M/M/1/25	9	10	0
M/D/1/25	3	4	0.9461
M/D/1/25	9	10	0.8581
M/U(0,0.5)/1/25	3	4	0.6358
M/U(0,0.2)/1/25	9	10	0.5791



**Figure 3.** Adjusted diffusion approximation.

## NUMERICAL EXPERIMENTS


In this section we test the applicability of the approximation to non-Gaussian environment and discuss the effects of the algorithm on the flow level.

### Accuracy of the Approximation

We calculated both the exact prices and the corresponding approximations, as described in the previous section, for simple arbitrarily selected queuing models (for which the exact prices could be found). Figure 3 shows the approximations of M/M/1/K, M/D/1/K, and M/U/1/K (where U stands for the uniform distribution), with the parameters shown in the Table 1 (parameters  $\delta$  are chosen “optimally” to illustrate the potential of the approximation). M/M/1/K is approximated almost exactly and the two other, *substantially non-Gaussian*, processes can be approximated with a high accuracy.

### Flow Level Considerations

The penalty function (see the second Section) arises from the price flow and controls the flow level behaviour of the network. Figure 4 shows the prices (bottom) and the losses (top) in the M/M/1/K model as a function of load. Clearly marks are generated at a significantly higher rate than losses occur. Still,



**Figure 4.** Flow level behaviour of the pricing in an M/M/1/K model.

the penalty function penalises the problem lighter than the one resulting from the virtual queue [12]. A low overall marking rate is especially advantageous when a packet passes through several resources. Also the gentle slope of the function is desirable knowing that large values of the derivative of the pricing function may compromise the stability in the presence of propagation delays [19]. Furthermore, the price signals are not occurring in bursts (except in case of a real congestion event), which allows the users to react lightly. This feature could be further emphasised by using several bits for the prices instead of marking probabilistically.

## DISCUSSION AND CONCLUSION

We proposed a packet marking algorithm based on system prices. The algorithm marks the known prices directly on the packets, while the unknown ones are replaced by their expected values. Calculation of the expected values depends on a suitable system model. We validated the algorithm for a simple M/M/1/K model, gave it a heuristic MDP interpretation, and finally generalised the approach for any non-correlated traffic.

The implementation of the congestion pricing scheme requires a simple, fair and robust marking algorithm. Our proposed algorithm is simple as the marking decisions are based directly on the state information, which requires only two parameters to be estimated from the traffic. It attempts also to be fair by pricing correctly on average, which is approximately achieved at least for non-correlated traffic. Finally, it is robust, since the deterministic component of the algorithm does not require any assumptions on the traffic.

In all, the algorithm provides a natural early warning feature comparable to the virtual queue, but with a direct approximate connection to the actual prices. Although some averaging is

required at the resources for parameter estimation, the method is still, by definition, more reactive than those directly based on the average queue size.

The limitations of the approach are related to parameter estimation and model selection. In parameter estimation we face the obvious trade-off when selecting the time we take into consideration in the estimation process. A longer time provides more accurate estimates but, under fluctuating traffic, it can increase bias in the expectation of the estimators and, on the other hand, make the model slower to react to changes. The correction term for non-Gaussian features may also be difficult to extract from the traffic, although one option would be to measure the overflow probability from the state  $K - 1$  and solve  $\delta$  from the expression (26) for  $u(K - 1 + \delta)$ .

As to the traffic model, we have assumed that the users have identical packet size distribution, and that the traffic is non-correlated. The significance of these undoubtedly restricting assumptions is unclear on time scales comparable to the busy periods and requires further work by means of simulation. However, it seems unlikely that the use of correlated traffic models would result in practical improvements in overall situation due to the increasing complexity in e.g. online estimation of parameters and correlation structures. Instead, we believe that future enhancements in the marking algorithms require additional available information on the flows using the resource.

## Acknowledgments

This work was funded by the Academy of Finland. Aleksi Penttinen received support also from the Nokia Foundation. The authors wish to thank Samuli Aalto for his contribution to the heuristic derivation of the MDP result, and Jouni Karvo for useful discussions.

## REFERENCES

- [1] Richard J. Gibbens and Frank P. Kelly, "Resource Pricing And the Evolution of Congestion Control," *Automatica*, vol. 35, pp. 1969–1985, 1999.
- [2] Frank P. Kelly, "Charging And Rate Control For Elastic Traffic," <http://www.statslab.cam.ac.uk/~frank/elastic.html>, 1997, Corrected version.
- [3] Frank P. Kelly, Amam K. Maulloo, and David K.H. Tan, "Rate Control in Communication Networks: Shadow Prices, Proportional Fairness and Stability," *Journal of the Operational Research Society*, vol. 49, pp. 237–252, 1998.
- [4] Frank P. Kelly, Peter B. Key, and Stan Zachary, "Distributed Admission Control," *IEEE Journal on Selected Areas in Communications*, vol. 18, no. 12, pp. 2617–2628, Dec. 2000.

- [5] Peter B. Key, Derek R. McAuley, Paul Barham, and Koenraad Laven, "Congestion Pricing for Congestion Avoidance," Tech. Rep. MSR-TR-99-15, Microsoft Research, Feb. 1999.
- [6] K. K. Ramakrishnan, Sally Floyd, and D. Black, "The Addition of Explicit Congestion Notification (ECN) to IP," <http://www.aciri.org/floyd/papers/draft-ietf-tsvwg-ecn-00.txt>, 2000.
- [7] Sally Floyd and Van Jacobson, "Random early detection gateways for congestion avoidance," *IEEE/ACM Transactions on Networking*, vol. 1, no. 4, pp. 397–413, Aug. 1993.
- [8] Sanjeewa Athuraliya and Steven H. Low, "Optimization Flow Control, II: Random Exponential Marking," Submitted for publication, <http://www.ee.mu.oz.au/staff/slow/research/projects.html>, May 2000.
- [9] Damon Wischik, "How to Mark Fairly," Workshop on Internet Service Quality Economics (MIT), 1999.
- [10] Frank P. Kelly, "Models for a Self-Managed Internet," *Philosophical Transactions of the Royal Society*, vol. A358, pp. 2335–2348, 2000.
- [11] Frank P. Kelly, "Mathematical Modelling of the Internet," in *Mathematics Unlimited - 2001 and Beyond*, B. Engquist and W. Schmid, Eds., pp. 685–702. Springer-Verlag, Berlin, 2001.
- [12] Richard Gibbens and Peter Key, "Distributed Control and Resource Marking Using Best-Effort Routers," *IEEE Network*, vol. 15, no. 3, pp. 54–59, May 2001.
- [13] Peter Key, "Resource Pricing for Differentiated Services," Available at <http://www.research.microsoft.com/users/pbk/>, Feb. 2001.
- [14] Lennart Råde and Bertil Westergren, *Mathematics Handbook for Science and Engineering*, Studentlitteratur, Lund, 3rd edition, 1995.
- [15] Alekski Penttinen, "Mathematical Models for Marking in Congestion Pricing," M.Sc. Thesis, Networking Laboratory, Helsinki University of Technology, Aug. 2001.
- [16] K. R. Krishnan, "Markov Decision Algorithms for Dynamic Routing," *IEEE Communications Magazine*, pp. 66–69, Oct. 1990.
- [17] Samuel Karlin and Howard M. Taylor, *A Second Course in Stochastic Processes*, Academic Press, San Diego, 1981.
- [18] Peter G. Harrison and Naresh M. Patel, *Performance Modelling of Communication Networks and Computer Architectures*, Addison-Wesley, Wokingham, 1992.
- [19] Frank P. Kelly, "Congestion Control: Fairness, Pricing And Stability," <http://www.statslab.cam.ac.uk/~frank/TALKS/ccfps.html>, 2000.

**Alekski Penttinen** received the M.Sc. (Tech) in engineering physics from the Helsinki University of Technology in 2001. He is currently a D.Sc. student at the Networking Laboratory of the Helsinki University of Technology with ad hoc networking as a special area of interest.

**Jorma Virtamo** received the M.Sc. (Tech) degree in engineering physics and D.Sc. (Tech) degree in theoretical physics from the Helsinki University of Technology, in 1970 and 1976, respectively.

In 1986, he joined the Technical Research Centre of Finland, VTT Information Technology, where he led a teletraffic research group, and became a Research Professor in 1995. Since 1997 he has been a Professor in the Networking Laboratory of the Helsinki University of Technology. His current research interests include queueing theory and performance analysis and control of the Internet, optical networks and ad hoc networks.