Alcom-FT Technical Report Series
ALCOMFT-TR-01-64

Latency Lower Bounds for Randomized Dynamic Bandwidth Allocation*

J. Garofalakis [1,2]   S. Nikoletseas [1,2]   P. Spirakis [1,2]

(1)
Computer Technology Institute (CTI), Patras, Greece
emails: {garofala, nikole, spirakis}@cti.gr
contact author: Paul Spirakis

(2)
Computer Engineering and Informatics Department,
Patras University, Greece.

Abstract

In this work we study the limitations of on-line schemes used to deal with the interplay of latency and utilization parameters in bandwidth allocation for tasks with varying and unpredictable bandwidth demands. We provide tight latency lower bounds to any randomized adaptive bandwidth allocation scheme for a single worst-case input stream of packets (session). We achieve this by employing the min-max principle of Yao and by characterizing the optimal (with respect to average queue size) deterministic scheme, which is subjected to a suitably chosen ("worst case") random input.

*This work was partially supported by the EU Projects ESPRIT LTR ALCOM- IT, IST-FET-OPEN ALCOM-FT , IMPROVING RTN ARACNE and the Greek GSRT Project PENED-ALKAD.

1 Introduction

Current (and future) communication networks are inviting an inflation in their bandwidth demand, due to users request for *quality of service* (QoS). Although the available bandwidth is increasing dramatically, it is still one of the bottleneck resources in such networks. Quality of service is characterized by many parameters. Among the most important are *latency (delay) and utilization*. (Examples of this trend, which allows some of the requirements for QoS to be specified explicitly by the user, can be found in the ATM standards and in the IPv6 standard).

When considering bandwidth allocation for tasks with varying and unpredictable bandwidth requirements, the latency and utilization parameters interplay. To understand this, remark that if the latency is measured by size of queue, then an overallocation of bandwidth may zero the queue size most of the time but allocate a lot of unused bandwidth (rate of service). Reverse, an underallocation of bandwidth might increase the queue lengths in an unbounded way. Clearly, latency is one of the most important parameters from the end user perspective (as it can be easily related to by any "web surfer"), but utilization on the other hand, is one of the most important performance parameters from the network point of view.

Amotz Bar-Noy et al ([1]) proposed models for dynamic bandwidth allocation and provided online algorithms that minimize *the number of bandwidth allocation changes* (this is technically important since it reduces the switch overhead which requires the invocation of special software in any such change). For other trends in applying dynamic bandwidth allocation see [5], [4] and [2].

In this work, we consider a single input session (stream of equisized small packets) of varying and unpredictable *arrival rate* (bandwidth requirement). We model any on-line allocation scheme by *a switch* with a single buffer (queue) which can dynamically adjust its *service rate* (allocated bandwidth) according to queue size. Our main goal is to understand the limitations of such adaptive schemes, i.e. to provide lower bounds for a whole class of allocation algorithms. To do this, we carefully employ Yao's min-max principle (see e.g. [9, 10]) which relates lower bounds of any (randomized) scheme for worst case inputs with (lower bounds to) the performance of an *optimal* deterministic algorithm for suitably chosen *random* inputs. We chose the Poisson stream as the random input (whose rate, λ, will be suitably selected to give the "worst case" behaviour leading to tight bounds when applying the min-max principle) and the straightforward deterministic adaptive allocation scheme which forces the "switch" to serve with a new rate when the queue size changes between certain values.

Although we use the problem of discovering on "optimal" deterministic scheme as a step in applying the min-max principle to derive the lower bounds in mind, we also note here that the problem of optimizing the parameters of such a deterministic adaptive scheme is an important, open till now, problem in *non-linear* stochastic optimization. The problem becomes even more interesting and technically difficult to deal with due to the realistic assumption that the available bandwidth rates of the algorithm are *integers*.

2 The adaptive bandwidth allocation switch

We consider here a switch which has an unbounded buffer. Packets arrive at the switch, creating a queue, and the switch serves them at a rate of packets/sec which is adaptively decided. The arrival process is worst-case. We seek to optimize the on-line algorithm by which the switch serves (with respect to its expected steady-state queue length), under the constraint that the switch's utilization remains always high.

2.1 The model

We consider the following system: the input traffic is modelled as any choice of a set I of a finite number of Poisson distributions. Let P_λ be such a choice, characterized completely by its rate of arrivals λ.

The switch (bandwidth allocation) algorithms that we consider is a finite set $\mathcal{A}(q_{\max})$ of deterministic algorithms which work as follows:

When the number of packets in the queue of the switch is more than a specified parameter $q_{\max}$, then the switch serves the queue with service rate

$$\mu_{\max} = \frac{1}{s_{\max}}$$

i.e. with service time $s_{\max}$ per packet.

Now, each algorithm $A \in \mathcal{A}$ chooses a suitable integer $k \leq q_{\max}$ and boundaries for the queue size

$$0 < q_1 < q_2 < \cdots < q_{\max} = q_{k-1}$$

(k intervals) and allocates to each such interval a service rate

$$\mu_i = \frac{1}{s_i}$$

(where s_i, $1 \leq i \leq k$, are the corresponding switch service times), so that, when the number of packets in the queue lies between q_{i-1} and q_i then the service rate of the switch is μ_i. Also, we demand that

$$\forall i, \ |q_i - q_{i-1}| \geq \Delta_q$$

where Δ_q is a given bandwidth restriction parameter. We assume that the (mean) service times s_i are all integers and of course $s_{\max} = s_1 \geq s_2 \geq \cdots \geq s_k$. Let the maximum of the Poisson arrival rates be

$$\lambda_{\max} = \max\{\lambda : P_\lambda \in I\}$$

For stability reasons then

$$\lambda_{\max} < \mu_{\max}$$

Note that $|\mathcal{A}|$ is then finite.

Clearly, any choice of an algorithm $A \in \mathcal{A}$ with input P_λ leads to a (mean) number $N(A, P_\lambda)$ of packets in the queue (including the packet in service) and a utilization $\rho(A, P_\lambda)$ in the steady-state operation of the switch. The *optimal* (with respect to queue size) deterministic algorithm A_{opt} for the (arbitrarily chosen) input P_λ is the algorithm A that minimizes $N(A, P_\lambda)$ with respect to the high utilization constraint

$$\rho(A, P_\lambda) \geq 1 - \epsilon$$

(where $\epsilon > 0$ given and small).

By Yao's minimax principle then,

$$N(A_{opt}, P_\lambda) \leq \max_{P_\lambda \in I} N(A_q, P_\lambda)$$

where A_q is *any randomized* algorithm serving the switch's queue adaptively (on-line), for any distribution q.

Our work then seeks (a) in characterizing the optimal deterministic algorithm A_{opt} for given input P_λ and (b) in providing (by determining a "worst case" random input when applying Yao's principle) a tight lower bound to any randomized technique which adapts to the requirements of the arrival pattern.

Note that the finiteness of I and $\mathcal{A}$ are only needed for the application of Yao's minimax principle. The integer stochastic non-linear optimization problem of minimizing the expected queue length of the switch given a lower bound on its utilization is an important problem of independent interest.

Notice that, to prove a lower bound on randomized complexity (expected queue size), it suffices to choose some distribution on the input and prove a *lower bound* on the expected queue size of deterministic algorithms for that distribution. It is important to remember that the deterministic algorithm "knows" the chosen input distribution.

Notice that the difficulty here is to prove a lower bound on the expected queue length of *any* deterministic algorithm on the chosen input. To solve this problem, we first construct (in an efficient way) the *optimal* deterministic algorithm and then obtain a (tight) lower bound for its mean queue size.

2.2 The problem formulation

We are given a set of allowable switch service times

$$S = \{s_1, s_2, \ldots, s_t = s_{\max}\}$$

where

$$s_1 \leq s_2 \leq \cdots \leq s_t = s_{\max}$$

i.e. a corresponding set R of integer service rates

$$R = \{\mu_1, \mu_2, \ldots, \mu_t\}$$

where

$$\forall i, \quad \mu_i = \frac{1}{s_i}$$

thus

$$\mu_1 \geq \mu_2 \geq \cdots \geq \mu_t$$

and a bandwidth restriction $\Delta_q > 0$ (Δ_q integer) obliging the switch to keep each selected switch service time (and, correspondingly, rate) for at least Δ_q contiguous numbers of packets in the switch queue (including the packet in service). For simplicity we assume $\mu_t = 1$. We wish to separate the integer line $\{1, 2, \ldots, \infty\}$ corresponding to the number of packets in the system into (at most) k consecutive segments, of length at least Δ_q each, so that each segment is assigned a switch rate (of the values in R) indicating that the switch serves with that rate for all the segment of numbers of packets. We further wish to assign the rates in such a way so that if the input (arrival process) is Poisson arrivals of rate $\lambda > 0$ then the expected number of packets in the system (steady-state) *is minimized and also* the switch utilization ρ (the steady-state percentage of time during which the switch is serving) is at least a given value $1 - \epsilon$ ($\epsilon > 0$ small). Such a switch is optimal with respect to the mean queue length number under the high utilization restriction. Note that, independently of the use of an optimal switch in the Yao's lower bound technique, its discovery is an interesting problem in integer stochastic nonlinear optimization.

3 An efficient algorithm for the construction of the optimal bandwidth allocation solution

3.1 Steady-state system equations

Assume that the (unknown) segments of the number of packets in the system are

$$[0, q_1], [q_1 + 1, q_2], \ldots, [q_{k-1} + 1, \infty] \quad (k \leq t)$$

Let j be the switch population (number of packets in the queue). Then for $q_{i-1} < j \leq q_i$ the switch behaves as an $M|D|1$ queue (Poisson input, fixed service rate). Let the rate assigned to this interval be r_i, where all r_i's are different. Obviously, $\{r_1, \ldots, r_k\} \subseteq R$.

Definition 1 *The branch system* $\sum_i$ *is an* $M|D|1$ *queue of Poisson input rate* λ *and fixed service rate* r_i*. Let* N_i *be its expected system population, in the steady state.* □

Remark: By the well-known Pollaczek-Khinchin formula for $M|D|1$ (see [6], vol. I, p. 188)

$$N_i = \frac{\frac{\lambda}{r_i}}{1 - \left(\frac{\lambda}{r_i}\right)^2} - \frac{\left(\frac{\lambda}{r_i}\right)^2}{2\left(1 - \frac{\lambda}{r_i}\right)}$$

Definition 2 *Let* b_{i+1} $(i = 0, \ldots, k-1)$ *be the steady-state probability that the number of packets* j *in the switch (buffer and in-service) is such that* $q_i < j \leq q_{i+1}$. □

Note. b_i is the probability that the switch "behaves as" the branch system $\sum_i$ (note that each branch system $\sum_i$ is assumed to have a fixed service rate for any number of packets in the queue).

Definition 3 *Let* $p_i(j)$ *be the steady-state probability of* j *packets in the branch system* $\sum_i$. □

Consider the distribution

$$\forall j \geq 0, \tilde{p}(j) = b_1 \cdot p_1(j) + \cdots + b_k \cdot p_k(j)$$

Clearly, this is a probability distribution over $j \geq 0$ since

$$\sum_{j=0}^{\infty} \tilde{p}(j) = b_1 \sum_{j=0}^{\infty} p_1(j) + \cdots + b_k \sum_{j=0}^{\infty} p_k(j) = \sum_{i=0}^{k} b_i = 1$$

and each $\tilde{p}(j) \geq 0$.

Now,

$$\Pr\{q_{i-1} \leq j < q_i, \;\; w.r.t. \;\; \tilde{p}\} = \sum_{j=q_{i-1}}^{q_i - 1} \tilde{p}(j)$$

By setting this equal to b_i we get the following system of k (fixpoint) equations

$$\forall i = 0, \ldots, k-1, \quad b_i = b_1 \sum_{j=q_{i-1}}^{q_i - 1} p_1(j) + \cdots + b_k \sum_{j=q_{i-1}}^{q_i - 1} p_k(j) \tag{1}$$

(these equations are not independent of each other since their sum leads to $\sum_{i=0}^{k-1} b_i = 1$).

Any $k-1$ of these fixpoint equations, together with

$$\sum_{i=0}^{k-1} b_i = 1$$

leads to a unique set of values for $b_1, \ldots, b_k$, for fixed system parameters.

Note that the equations in (1) provide an approximate evaluation of the values of b_i. Note also that $\forall j, \tilde{p}(j)$ is an approximate estimation of the steady-state probability of j packets in the switch.

Definition 4 *Let $\tilde{S}$ be the system defined by $\tilde{p}()$.* □

Definition 5 *Let $\tilde{N} = \sum_1^k b_i N_i$.* □

Theorem 1 *$\tilde{N}$ is an upper bound on the expected number of packets in the switch S in the steady-state.*

Proof (sketch): Each branch system contributes to the other branches in $\tilde{p}()$, while this is not allowed in the considered switch S. But then the system $\tilde{S}$ contributes a greater variance (than S) to the steady-state. The Theorem then follows from the Pollaczek-Khinchin mean value formula (see e.g. [6], Vol. I, p. 187) for any $M|G|1$ system (for general service time distribution this formula gives

$$\bar{N} = \lambda \bar{x} + (\lambda \bar{x})^2 \frac{1 + \frac{\sigma_b^2}{\bar{x}^2}}{2(1 - \lambda \bar{x})}$$

where $\bar{x}$ is the expected service time of the switch and σ_b^2 its variance). □

Note. A similar phenomenon is happening when we compare the mean buffer length of an $M|M|1$ and an $M|D|1$ (deterministic constant service rate) queues of the same mean service rate. The system $M|D|1$ has a smaller queue length on the average (see [6], vol. I, p. 188), because of the fact that the expected queue length increases with the variance of the service-time distribution.

3.2 The algorithm for determining the optimal switch

We initially set the switch to the (maximum) rate μ_1 for a (high) number of packets in the system $q > q_{\max}$ and to $\mu = 1$ for $q \leq q_{\max}$. We then progressively adjust the switch rates in order to reduce the steady-state mean number of packets in the system by choosing the highest possible service rate and the longest possible interval of application of this rate while keeping the utilization high as follows:

Algorithm Opt
Initialization
$R_{current} := R - \{\mu_1\}$
$q_{current} := q_{\max}$

(a) Let $m := \max\{R_{current}\}$
(i.e. the maximum of the remaining available rates).
If $R_{current} = \emptyset$ then exit.
(b) Let $\Delta \geq \Delta_q$ be the biggest possible integer for which, if we set the switch rate to m for $q_{current} - \Delta \leq q \leq q_{current}$ (and to 1 for $[0, q_{currrent} - \Delta]$), we do not violate the utilization constraint.
If no such Δ exists, then

$$R_{current} := R_{current} - \{m\}$$

and go to step (a).
Else,

$$q_{currrent} := q_{currrent} - \Delta$$

and

$$R_{current} := R_{current} - \{m\}$$

and go to step (a).
Endof Opt

Note. To find the current utilization, $\rho_{current}$, needed in step (b), we calculate first the branch probabilities b_i and then

$$\rho_{current} := \sum \left(\frac{\lambda}{\mu_i}\right) b_i$$

where the sum is taken for the current set of intervals and rates.

3.3 The correctness of Algorithm Opt

Remark that each iteration of algorithm Opt progressively increases the rate of the switch in a new interval (where before we had a rate equal to 1) while keeping the high utilization constraint. Since the mean number N of packets in S at steady-state is a decreasing function of service rates (and their intervals of application), then N drops in each iteration until we cannot properly assign rates any more. Note that step (b) finds the longest interval of application of each rate (if this rate is admissible due to the utilization constraint). So we get:

Theorem 2 *Algorithm Opt determines the optimal (with respect to the mean number of packets in the system) switch S for Poisson input of arrival rate λ.* □

3.4 The efficiency of Algorithm Opt

In order to estimate the progress of algorithm Opt, we remark that the algorithm will terminate after a maximum of

$$O\left(\frac{q_{\max}}{\Delta_q}\right)$$

iterations. Let t be the (maximum) time needed to execute each iteration of algorithm Opt. Since in each iteration we calculate Δ by an implied binary search on the possible interval lengths and then solve the (at most $|R|$) system equations to approximately evaluate the branch probabilities b_i, we get

$$t = O(|R|^3 + \log q_{\max})$$

leading to a total execution time of algorithm Opt as in the following:

Theorem 3 *The execution time of Algorithm Opt is bounded above by*

$$O\left(|R|^3 \frac{q_{\max}}{\Delta_q} + \frac{q_{\max} \log q_{\max}}{\Delta_q}\right)$$

and thus it is polynomial in the system parameters. □

4 Lower bounds for any randomized bandwidth allocation scheme

Let $N, \tilde{p}(), \{b_i\}$ be respectively the mean number of packets in the queue, the probability distribution leading to $\tilde{N}$ and the branch probabilities of the optimal switch determined by our algorithm.

Note that the optimal switch, S_{opt}, is determined by a sequence of intervals $I_1, \ldots, I_k$ and their corresponding service rates $\mu(I_1), \ldots, \mu(I_k)$.

Definition 6 *The effective service time S_{eff} of S_{opt} is given by*

$$S_{eff} = \frac{1}{\mu_{eff}} = \sum_{I_j} b_j \cdot \frac{1}{\mu(I_j)}$$

□

Now, let us replace S_{opt} by a switch S_{eff} whose service time is always equal to S_{eff}.

The expected system size of S_{eff} is always a *lower bound* of N since the system corresponding to S_{eff} is an $M|D|1$ (variance zero).

Definition 7 *Let N_{eff} be the expected number of packets in the system S_{eff} in the steady-state.* □

Lemma 1

$$N_{eff}(\lambda) = \frac{\lambda\, S_{eff}}{1-(\lambda\, S_{eff})^2} - \frac{(\lambda\, S_{eff})^2}{2(1-\lambda\, S_{eff})}$$

Proof: By the well-known Pollaczek-Khinchin formula for $M|D|1$. □

Now, let $\lambda_0 \leq \mu_{\max}$ be the value of λ maximizing $N_{eff}(\lambda)$. Note that S_{eff} is itself a function of λ. To find λ_0 we can start from finding the root where the first derivative of $N_{eff}(\lambda)$ becomes zero. Then we have to choose any root at most $\mu_{\max}$ which makes the second derivative negative. If no such root exists, then $\mu_{\max}$ is the needed value of λ_0. For any particular instantiation of the input parameters we can write down the function $S_{eff}(\lambda)$ and do the above procedure. However a homogeneous solution (as a function of the input parameters) is difficult to get in a closed form.

Theorem 4 *$N_{eff}(\lambda_0)$ is a lower bound on the expected buffer size of any randomized bandwidth allocation algorithm for worst case inputs.*

Proof: By Yao's minimax principle, since $N_{eff}(\lambda_0)$ is a tight lower bound ("worst case" random input) to the expected buffer size of the optimal switch S_{opt}. □

5 Conclusion and further work

We have applied here queuing theoretic notions and a careful efficient algorithmic construction of the optimal switch solution via which we have obtained tight lower bounds to any randomized bandwidth allocation scheme with respect to a single input stream. We currently examine the limits of the extension of our method to a network of switches and multiple input streams of packets.

References

[1] A. Bar-Noy, Y. Mansour and B. Schieber, "Dynamic Bandwidth Allocation", ACM PODC 98, pp. 31–39, 1998.

[2] Y. Afek, M. Cohen, E. Haalman and Y. Mansour, "Dynamic bandwidth allocation", In Proc. IEEE INFOCOM, pp. 880-887, 1996.

[3] A. Bar-Noy, R. Canetti, S. Kutten, Y. Mansour and B. Schieber, "Bandwidth allocation with preemption", in Proc. 25th ACM Symp. On Theory of Computing (STOC), pp. 616-625, 1995.

[4] I. Cidon, R. Rom and Y. Shavit, "Analysis of one way reservation algorithms", In Proc. IEEE INFOCOM, pp. 1256–1263, 1995.

[5] M. Grossglauser, S. Keshav and D. Tse, "RCBR: A simple efficient service for multiple time-scale traffic", in Proc. ACM SIGCOM, pp. 219–230, 1995.

[6] L. Kleinrock, "Queuing Systems", Vol I, Wiley eds.

[7] F. Pollaczek, "Uber eine Aufgabe der Wahrscheinlichkeitstheorie", I-II Math. Zeitschrift, 32, 64–100, pp. 729–750, 1930.

[8] A.V. Khinchin, "Mathematical Theory of Stationary Queues", in Mat. Sbornik 39, 73-84, 1932.

[9] R. Motwani and P. Raghavan, "Randomized Algorithms", Cambridge University Press, 1995.

[10] A. Yao, "Probabilistic Computations: Towards a Unified Measure of Complexity", in Proc. of the 17th Symp. on Foundations of Computer Science (FOCS), pp. 222-227, 1977.