COST279TD(03)03

Providing Differentiated Services by Load Balancing and Scheduling in MPLS Networks

Riikka Susitaival and Samuli Aalto [1]

Networking Laboratory, Helsinki University of Technology

Abstract

The capability of MPLS of explicit routing as well as of splitting of the traffic on several paths allows load balancing. Using known load balancing algorithms as a starting point, the main goal of this paper is to formulate optimization problems that differentiate service classes in terms of mean delay. The differentiation is achieved by using both WFQ-scheduling and routing. We have approximated the upper and lower bound of ratio of mean delays between classes achieved by work-conserving scheduling. The load balancing algorithm and the algorithms to differentiate classes are implemented and tested. As a result it is found that the level of service differentiation achieved by scheduling is quite low. The use of routing provides more flexibility, but at the expense of longer mean delays.

1 INTRODUCTION

In the conventional IP networks, forwarding decisions are made independently in each router, based on the packet's header and precalculated routing tables. MPLS (Multi Protocol Label Switching) [12] is a flexible technology that enables new services in IP networks and makes forwarding more effective. MPLS is based on short fixed length labels, which are assigned to each packet at the ingress node of the MPLS cloud. These labels are used to make forwarding decisions at each node. The assignments of a particular packet to a particular FEC (Forwarding Equivalence Class) are made only once. This simplifies forwarding and makes it more effective.

One of the most significant applications of MPLS is Traffic Engineering (TE) [1]. TE using MPLS provides mechanisms to route traffic with the same origin and destination along several paths. The most important benefit of this kind of traffic splitting is the ability to balance the load. The load balancing reduces congestion and therefore improves the performance of the network. Assuming that the network topology, link capacities and traffic demands are known load balancing can be formulated as an optimization problem in which the mean packet delay is minimized, see e.g. [2].

The conventional IP networks offer services only on the best-effort basis. However, the demand to provide different services to different customers is evident. Quality of Service (QoS) is under a widespread discussion today. Two notable architectures, Integrated Services [4] and Differentiated Services [3], have been introduced in the literature. MPLS and its TE capabilities provide technical support to the implementation of QoS. The service differentiation can be obtained, to some extent, by a suitable flow allocation. The differentiation can be boosted by appropriate scheduling mechanisms, like WFQ-scheduling.

[1]E-mail: {riikka.susitaival, samuli.aalto }@hut.fi

Our goal is to formulate the service differentiation problem among two service classes as an optimization problem (similar to the load balancing problem), to solve the problem, apply the solution to a test network, and make some interesting observations from this application. We introduce three different formulations for the service differentiation problem. The first approach relies on the minimum delay routing applied to both classes. The service differentiation is then achieved by appropriate scheduling methods at the nodes. The second approach, on the contrary, uses classwise flow allocation without any specific scheduling mechanism (but just FIFO). The third approach tries to catch both the effects of classwise routing and scheduling mechanisms.

We emphasize that, due to lack of exact formulas for other than simple M/M/1-FIFO networks, we are obliged to make several (rather crude) approximations in our formulations. Thus, this study should be viewed more as an engineering approach than a strictly mathematical discipline. However, we think that the results are interesting enough to justify the approach. Also in [13] we have studied the same service differentiation problem, but the approach was based on the approximation of GPS by parallel queues.

The rest of this paper is organized as follows. Section 2 introduces the notation and formulation of the load balancing problem. In section 3 three approaches to differentiate service classes, differentiation by scheduling, differentiation by routing and differentiation by both scheduling and routing are presented. In the fourth section we analyze these methods by numerical examples. Finally, section 5 summarizes the paper.

2 LOAD BALANCING PROBLEM

In this section we formulate the load balancing problem as an optimization problem in which the mean packet delay is minimized. Here we follow the approach presented by Ott et al. [10, 5] rather than the original one [6]. This section serves as a preliminary for the following section, where the service differentiation problem is considered in the DiffServ context.

Consider a connected network $(\mathcal{N}, \mathcal{L})$ consisting of a set of nodes $\mathcal{N}$ and a set of unidirectional links $\mathcal{L}$. Denote the sizes of sets $\mathcal{N}$ and $\mathcal{L}$ by N and L, respectively. Each link $l \in \mathcal{L}$ is characterized by capacity c_l (bit/s). The network is loaded by traffic streams between any two nodes i and j. Let d_{ij} denote the average traffic rate (bit/s) for the stream (i, j) entering the network at node i and destined for node j. We assume that these traffic demands are known. The problem is to carry these demands optimally over the given network. This kind of load balancing problem can be formulated as follows.

Let $x_{ij,l}$ denote the part (bit/s) of traffic demand d_{ij} carried by link l. These flows $x_{ij,l}$ must satisfy the following *capacity constraints*:

$$f_l < c_l, \qquad \text{for each } l \in \mathcal{L}, \tag{1}$$

where f_l denotes the total average traffic rate carried by link l,

$$f_l = \sum_{i,j \in \mathcal{N}} x_{ij,l}.$$

In addition, these flows must satisfy the following *flow conservation equation*:

$$\mathbf{A}\mathbf{x}_{ij} = \mathbf{d}_{ij}, \qquad \text{for all } i, j \in \mathcal{N}, \tag{2}$$

where $\mathbf{x}_{ij} \in \mathbb{R}^{L \times 1}$ refers to the vector

$$[\mathbf{x}_{ij}]_l = x_{ij,l},$$

$\mathbf{d}_{ij} \in \mathbb{R}^{N \times 1}$ to the vector

$$[\mathbf{d}_{ij}]_n = \begin{cases} d_{ij}, & \text{if } n = i, \\ -d_{ij}, & \text{if } n = j, \\ 0, & \text{otherwise,} \end{cases}$$

and $\mathbf{A} \in \mathbb{R}^{N \times L}$ to the matrix

$$[\mathbf{A}]_{n,l} = \begin{cases} 1, & \text{if node } n \text{ is the origin of link } l, \\ -1, & \text{if node } n \text{ is the destination of link } l, \\ 0, & \text{otherwise.} \end{cases}$$

In addition to these constraints, we have to specify the objective function for the optimization. A reasonable one is derived from the Jackson queueing network model. Let X denote the total number of packets in the network. Assuming that all streams arriving in the network are Poissonian and all packet lengths are independent and exponentially distributed, the mean number of packets in the queue related to link l is

$$E[X_l] = \frac{f_l}{c_l - f_l}.$$

Thus, the total mean number of packets is

$$E[X] = \sum_{l \in \mathcal{L}} E[X_l] = \sum_{l \in \mathcal{L}} \frac{f_l}{c_l - f_l}. \quad (3)$$

By Little's result, this is proportional to the total mean delay $E[D]$ in the network,

$$E[D] = \frac{E[X]}{\lambda}$$

where λ refers to the total packet arrival rate to the network. Thus, minimizing the mean number of packets $E[X]$ is equivalent to minimizing the mean delay $E[D]$.

Load Balancing Problem (LB) *Find nonnegative $x_{ij,l}$'s that minimize the objective function* (3) *and satisfy the constraints* (1) *and* (2).

We note that the solution of the load balancing problem gives the linkwise portions $x_{ij,l}$ of the original demands d_{ij}, but the routes for these demands through the network are not necessarily unique. However, the routes are guaranteed to be loop-free.

3 SERVICE DIFFERENTIATION PROBLEM

In this section we consider the DiffServ context with the aim of achieving service differentiation among different classes $k \in \mathcal{K}$. More precisely, assume that there are just two service classes, $\mathcal{K} = \{1, 2\}$, and let the service differentiation be described by a single parameter, the ratio r of mean packet delays of the two classes,

$$r = \frac{E[D_2]}{E[D_1]}.$$

The target delay ratio is denoted by r^*. Thus, in this study we restrict ourselves to packet level delay differentiation. This might be a relevant approach to differentiate between EF and AF classes, or between two AF classes.

Consider a similar connected network $(\mathcal{N}, \mathcal{L})$ as in the previous section. The network is loaded by classwise traffic streams between any two nodes i and j. Let d_{kij} denote the average traffic rate (bit/s) for the stream (i, j) of class k. In addition, let d_k and λ_k denote the total average traffic rate (bit/s) of class k and the corresponding average packet arrival rate, respectively,

$$d_k = \sum_{i,j \in \mathcal{N}} d_{kij} \qquad \text{and} \qquad \lambda_k = \frac{d_k}{L},$$

where L refers to the mean packet size (in bits) in both classes. Now the problem is to carry the demands d_{kij} over the given network in such a way that the required service differentiation is achieved in an optimal way. After presenting below the necessary and sufficient conditions for the feasibility of the solution, we explore different approaches how to define this "optimality".

Let $x_{kij,l}$ denote the part (bit/s) of traffic demand d_{kij} carried on link l. These flows $x_{kij,l}$ must satisfy the following *capacity constraints*:

$$\sum_{k \in \mathcal{K}} f_{k,l} < c_l, \qquad \text{for each } l \in \mathcal{L}, \tag{4}$$

where $f_{k,l}$ denotes the average traffic rate of class k carried on link l,

$$f_{k,l} = \sum_{i,j \in \mathcal{N}} x_{kij,l}.$$

In addition, these flows must satisfy the following *flow conservation equation*:

$$\mathbf{A}\mathbf{x}_{kij} = \mathbf{d}_{kij}, \qquad \text{for all } k \in \mathcal{K} \text{ and } i,j \in \mathcal{N}, \tag{5}$$

where $\mathbf{x}_{kij} \in \mathbb{R}^{L \times 1}$ refers to vector

$$[\mathbf{x}_{kij}]_l = x_{kij,l}$$

and $\mathbf{d}_{kij} \in \mathbb{R}^{N \times 1}$ to vector

$$[\mathbf{d}_{kij}]_n = \begin{cases} d_{kij}, & \text{if } n = i, \\ -d_{kij}, & \text{if } n = j, \\ 0, & \text{otherwise.} \end{cases}$$

In addition to these constraints, we have to specify the objective function for the optimization. Below we present three different approaches to this problem.

3.1 First approach: service differentiation by scheduling

The first approach relies on the minimum delay routing applied to both classes. Thus, we assume that the aggregate flows $x_{ij,l}$ are determined as the solution of the Load Balancing Problem (LB) with aggregate demands

$$d_{ij} = \sum_{k \in \mathcal{K}} d_{kij}.$$

Clearly, this solution satisfies the capacity constraints (4). We split the aggregate flows $x_{ij,l}$ into classwise flows $x_{kij,l}$ by using the same proportions as in the network entrance,

$$x_{kij,l} = \frac{d_{kij}}{d_{ij}} x_{ij,l}. \qquad \text{for all } k \in \mathcal{K},\ i,j \in \mathcal{N},\ l \in \mathcal{L}. \tag{6}$$

By this way, the flow conservation equation (5) will be satisfied, as it is easy to see. Note, however, that this is not necessarily the only possible splitting that satisfies (5).

Service differentiation is now enabled, at least to a certain extent, by replacing the plain FIFO queues with appropriate scheduling methods. Let us try to find out to what extent.

One scheduling method often mentioned in the DiffServ context is Weighted Fair Queueing (WFQ). We will consider Generalized Processor Sharing (GPS) [11], which is the ideal version of WFQ. GPS applied in the output queue of link l is characterized by giving positive weights $\phi_{k,l}$ for each class k such that

$$\sum_{k \in \mathcal{K}} \phi_{k,l} = 1. \tag{7}$$

The link capacity c_l is divided proportional to these weights among the active classes, i.e. such classes that have packets in the buffer. We would like to know the resulting mean delay ratio r whenever the weights $\phi_{1,l}$ and $\phi_{2,l}$ for all the links l are given and the routing of both classes follows the minimum delay routing as explained above.

The problem with GPS (as well as with any WFQ-variant) is that it is not analytically tractable. Even if one makes the simplifying memoryless assumptions, i.e. the arrival processes of both classes are Poissonian and all the packet lengths are exponentially distributed, there is no analytical expression for the classwise mean packet delays. To tackle this problem, we use a numerical approximation based on the observation that a GPS queue with two classes approaches a pure priority queue (with pre-emption) as one of the weights approaches one. We start by determining approximative upper and lower bounds for the mean delay ratio r.

Approximative Upper Bound Assume first that all the FIFO queues are replaced by pre-emptive priority queues with class 1 having the highest priority (i.e., $\phi_{1,l} \equiv 1$ and $\phi_{2,l} \equiv 0$). By this way the mean delay ratio r clearly reaches its upper bound among all possible weight combinations $\phi_{1,l}$ and $\phi_{2,l}$.

By approximating the flows inside the network by corresponding Poissonian flows, we get the following approximations for the mean number of packets in the two classes related to link l (see e.g. [9]):

$$E[X_{1,l}] \approx \frac{f_{1,l}}{c_l - f_{1,l}} \qquad \text{and} \qquad E[X_{2,l}] \approx \frac{f_{2,l} c_l}{(c_l - f_{1,l})(c_l - f_{1,l} - f_{2,l})}.$$

Thus, by Little's result, we get the following approximative upper bound for the mean delay ratio r:

$$r_U \approx \frac{\frac{1}{\lambda_2} \sum_{l \in \mathcal{L}} \frac{f_{2,l} c_l}{(c_l - f_{1,l})(c_l - f_{1,l} - f_{2,l})}}{\frac{1}{\lambda_1} \sum_{l \in \mathcal{L}} \frac{f_{1,l}}{c_l - f_{1,l}}}, \tag{8}$$

where λ_k refers to the total packet arrival rate of class k.

Approximative Lower Bound Assume then that all the FIFO queues are replaced by pre-emptive priority queues with class 2 having the highest priority (i.e., $\phi_{1,l} \equiv 0$ and $\phi_{2,l} \equiv 1$). By this way the mean delay ratio r clearly reaches its lower bound, for which we have the following approximation (determined in a similar way as for the upper bound):

$$r_L \approx \frac{\frac{1}{\lambda_2} \sum_{l \in \mathcal{L}} \frac{f_{2,l}}{c_l - f_{2,l}}}{\frac{1}{\lambda_1} \sum_{l \in \mathcal{L}} \frac{f_{1,l} c_l}{(c_l - f_{2,l})(c_l - f_{1,l} - f_{2,l})}}. \tag{9}$$

Approximative GPS-Weights Given ratio requirement $r^* \in [r_L, r_U]$ a simple linear interpolation gives the following GPS-weights to be used in each queue l:

$$\phi_{1,l} = \frac{r^* - r_L}{r_U - r_L} \qquad \text{and} \qquad \phi_{2,l} = \frac{r_U - r^*}{r_U - r_L}. \tag{10}$$

So, this first approach gives us an approximative method to determine the required GPS-weights ϕ_{kl} assuming that our target ratio r^* falls into interval $[r_L, r_U]$. To conclude, let us formulate this as an optimization problem.

Service Differentiation Problem, Formulation 1 (SD-1) *Find nonnegative $x_{ij,l}$'s that minimize the objective function* (3) *and satisfy the constraints* (1) *and* (2). *Split these aggregate flows $x_{ij,l}$ into classwise flows $x_{kij,l}$ according to* (6). *Given the required service differentiation ratio $r^* \in [r_L, r_U]$, determine the GPS-weights $\phi_{k,l}$ from* (10).

The advantage of this method is its efficiency: since both classes follow the optimal mean delay routes, the service differentiation is achieved without any impairment in the total mean delay.

3.2 Second approach: service differentiation by routing

Contrary to the first approach, the second one adheres to the FIFO scheduling and tries to achieve the required service differentiation solely by routing methods. Below we formulate this service differentiation problem as an optimization problem in two slightly different ways.

Let us start with a feasible set of classwise flows $x_{kij,l}$, i.e., such that satisfy the constraints (4) and (5) mentioned above. Due to the latter constraint, end-to-end routes for all classwise demands d_{kij} can be determined from these classwise flows (while not necessarily uniquely). Assuming that the packets arrive to the network according to Poisson processes and that all the packet lengths are exponentially distributed, we have an open queueing network (of Kelly type [8]), for which the classwise total mean delays are as follows:

$$E[D_1] = \frac{1}{\lambda_1}\sum_{l\in\mathcal{L}}\frac{f_{1,l}}{c_l - f_{1,l} - f_{2,l}} \qquad \text{and} \qquad E[D_2] = \frac{1}{\lambda_2}\sum_{l\in\mathcal{L}}\frac{f_{2,l}}{c_l - f_{1,l} - f_{2,l}}. \tag{11}$$

In the first formulation we pose a new constraint that enforces the required service differentiation ratio:

$$r = r^*, \tag{12}$$

where

$$r = \frac{E[D_2]}{E[D_1]} = \frac{\frac{1}{\lambda_2}\sum_{l\in\mathcal{L}}\frac{f_{2,l}}{c_l - f_{1,l} - f_{2,l}}}{\frac{1}{\lambda_1}\sum_{l\in\mathcal{L}}\frac{f_{1,l}}{c_l - f_{1,l} - f_{2,l}}}. \tag{13}$$

To improve efficiency, we are still free to optimize one of the classwise delays. This leads to the following optimization problem.

Service Differentiation Problem, Formulation 2a (SD-2a) *Find nonnegative $x_{kij,l}$'s that minimize the mean delay $E[D_1]$ in* (11) *and satisfy the constraints* (4)*,* (5) *and* (12)*.*

We have observed that practical methods to solve this optimization problem may generate loops. The problem is alleviated (maybe even solved) by first relaxing the ratio requirement (12) and introducing classwise cost weights $w_k > 0$. So our objective function is now

$$\sum_{k\in\mathcal{K}} w_k E[D_k] = \sum_{l\in\mathcal{L}}\frac{\frac{w_1}{\lambda_1}f_{1,l} + \frac{w_2}{\lambda_2}f_{2,l}}{c_l - f_{1,l} - f_{2,l}}, \tag{14}$$

and we have the following optimization problem.

Service Differentiation Problem, Formulation 2b (SD-2b) *Find nonnegative $x_{kij,l}$'s that minimize the objective function* (14) *and satisfy the constraints* (4) *and* (5)*.*

As the result of this optimization we get the classwise flows $x_{kij,l}$. Under the same assumptions as before, they result in the delay ratio r calculated from (13). To achieve the required ratio r^* this optimization procedure has to be iterated.

3.3 Third approach: service differentiation by scheduling and routing

Our third approach has, in principle, the highest degree of freedom, since we allow both routing and scheduling methods to be used to achieve the required service differentiation. However, we restrict ourselves here just to extending the first approach beyond the limits r_L and r_U. For example, if the required delay ratio r^* exceeds the upper bound r_U, it is then natural to use pure priority scheduling (with class 1 having the highest priority) and give up the minimum delay routing to allow classwise routes to achieve the target delay ratio. Another possibility would be to formulate the most general optimization problem, which takes both the classwise flows $x_{kij,l}$ and GPS-weights $\phi_{k,l}$ as free parameters. However, in that case, we should have a good approximation for classwise delays in a GPS-scheduled queue.

Assume first that $r^* > r_U$. Consider now the case where each link is provided with a pure priority output queue that gives the highest priority to class 1 and the network is loaded by a feasible set of classwise flows $x_{kij,l}$. As in subsection 3.1, we approximate the flows inside the network by corresponding Poissonian flows, which leads to the following formulae for the total mean delay in the two classes:

$$E[D_1] \approx \frac{1}{\lambda_1} \sum_{l \in L} \frac{f_{1,l}}{c_{1,l} - f_{1,l}} \qquad \text{and} \qquad E[D_2] \approx \frac{1}{\lambda_2} \sum_{l \in L} \frac{f_{2,l} c_l}{(c_l - f_{1,l})(c_l - f_{1,l} - f_{2,l})}. \tag{15}$$

From this on, we follow the lines presented in the previous subsection. So, we pose a new constraint that enforces the required service differentiation ratio:

$$r = r^* > r_U, \tag{16}$$

where

$$r = \frac{E[D_2]}{E[D_1]} \approx \frac{\frac{1}{\lambda_2} \sum_{l \in \mathcal{L}} \frac{f_{2,l} c_l}{(c_l - f_{1,l})(c_l - f_{1,l} - f_{2,l})}}{\frac{1}{\lambda_1} \sum_{l \in \mathcal{L}} \frac{f_{1,l}}{c_{1,l} - f_{1,l}}}. \tag{17}$$

To improve the efficiency, we are again free to optimize one of the classwise delays. This leads to the following optimization problem.

Service Differentiation Problem, Formulation 3a (SD-3a) *Find nonnegative $x_{kij,l}$'s that minimize the mean delay $E[D_1]$ in* (15) *and satisfy the constraints* (4), (5) *and* (16).

Assume then that $r^* < r_L$ and consider the case where each link is provided with a pure priority output queue that gives the highest priority to class 2. Along similar lines as above, any feasible set of classwise flows $x_{kij,l}$ leads to the following formulas for the total mean delay in the two classes:

$$E[D_1] \approx \frac{1}{\lambda_1} \sum_{l \in L} \frac{f_{1,l} c_l}{(c_l - f_{2,l})(c_l - f_{1,l} - f_{2,l})} \qquad \text{and} \qquad E[D_2] \approx \frac{1}{\lambda_2} \sum_{l \in L} \frac{f_{2,l}}{c_{1,l} - f_{2,l}} \tag{18}$$

From this on, we follow the lines presented in the previous subsection. So, we pose a new constraint that enforces the required service differentiation ratio:

$$r = r^* < r_L, \tag{19}$$

where

$$r = \frac{E[D_2]}{E[D_1]} \approx \frac{\frac{1}{\lambda_2} \sum_{l \in \mathcal{L}} \frac{f_{2,l}}{c_{1,l} - f_{2,l}}}{\frac{1}{\lambda_1} \sum_{l \in \mathcal{L}} \frac{f_{1,l} c_l}{(c_l - f_{2,l})(c_l - f_{1,l} - f_{2,l})}}. \tag{20}$$

To improve efficiency, we are again free to optimize one of the classwise delays. This leads to the following optimization problem.

Service Differentiation Problem, Formulation 3b (SD-3b) *Find nonnegative $x_{kij,l}$'s that minimize the mean delay $E[D_1]$ in* (18) *and satisfy the constraints* (4), (5) *and* (19).

Remark: to avoid routing loops, formulations (SD-3a) and (SD-3b) may be replaced by a similar weighted mean delay approach as presented in the previous subsection.

We note that, due to crude approximations made above, it is difficult to fairly compare the proposed solutions to achieve the required level of service differentiation without simulations. This will be a topic of future research.

4 NUMERICAL RESULTS

The optimization problems LB, SD-2a, SD-2b, SD-3a and SD-3b described above are formulated using a General Algebraic Modelling System (GAMS) [7], which is a high-level modelling language for mathematical programming specially suited for large scale optimization problems. After the formulation in GAMS the problems are solved using Minos 5 solver module.

We have tested optimization algorithms in a test-network, which consists of 10 nodes, 52 unidirectional links and 72 traffic streams. The topology $(\mathcal{N}, \mathcal{L})$, the link capacities c_l and the aggregate traffic demands d_{ij} of the network are taken from web-page http://brookfield.ans.net/omp/random-test-cases.html.

4.1 Load balancing problem

First we consider the load balancing problem (LB), which excludes the service classes. The mean delay achieved by the routing with load balancing is compared to the shortest path routing, where the traffic streams are routed along the paths that minimize the hop count. The mean delay as a function of the traffic load is presented in Figure 1.

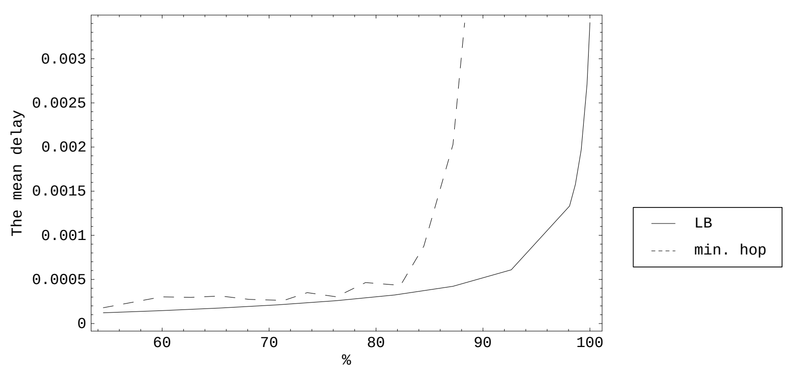

Figure 1: The mean delay as a function of the traffic load

As expected, the use of load balancing reduces the mean delay especially when the traffic load is heavy (over 80 % of the maximum traffic load).

4.2 Service differentiation by scheduling

Next we take the service classes into account. We assume that the aggregate traffic demand of each stream is divided between the service classes with fixed proportions. Let a_k denote the proportion of the traffic demand of class k of total aggregate traffic demand, that is,

$$\frac{d_{kij}}{d_{ij}} = a_k \qquad \text{for all } i, j \in \mathcal{N} \text{ and}$$

$$\sum_{k \in \mathcal{K}} a_k = 1.$$

The approximative lower (9) and upper (8) bounds for the ratio of mean delays, $E[D_2]/E[D_1]$, as a function of a_1 in our test-network are presented in Figure 2.

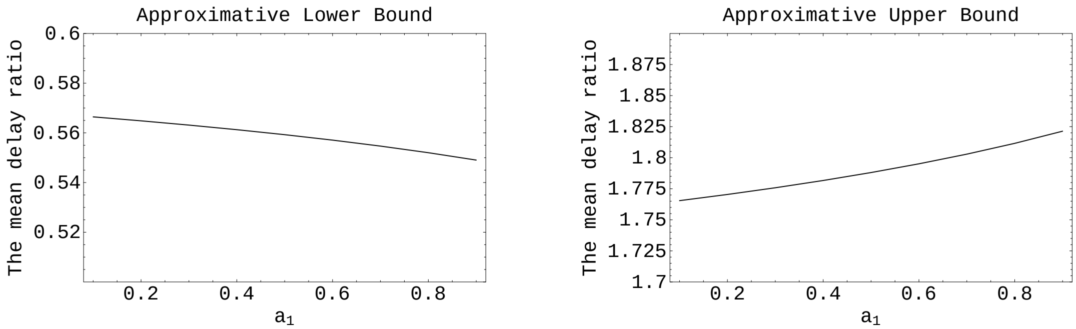

Figure 2: The approximative lower and upper bounds for the ratio of classwise mean delays as a function of a_1

Finding the classwise flow allocation by dividing the aggregate link traffic in accordance with equation (6) is not the only possibility. The performance of minimum-delay routing can be obtained by fixing the aggregate link flows on each link but restricting the link flows of each class only by (5). This approach provides a bit more flexibility to differentiate classes. If we optimize the mean delay of the first class, we may obtain the ratio of mean delays that is greater than the approximative upper bound presented in Figure 2. Figure 3 shows the upper bounds (8) as a function of a_1 by fixing the link flow of each class ("class") and by fixing only the aggregate link flow ("aggregate").

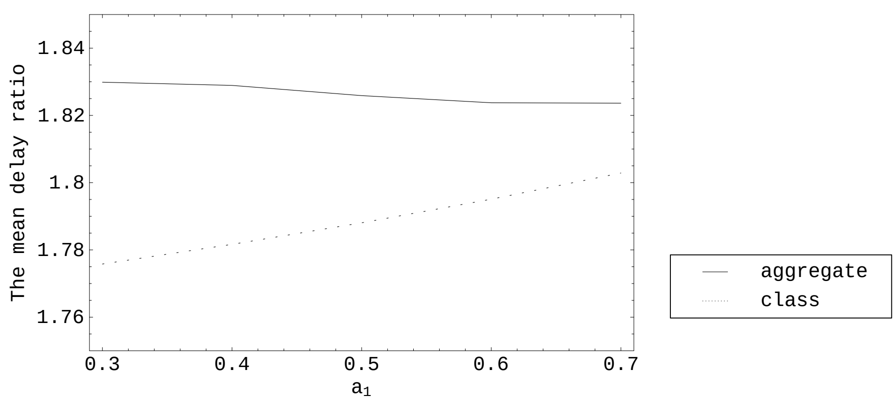

Figure 3: The approximative upper bound for the classwise ratio of mean delays as a function of a_1

4.3 Service differentiation by routing

The objective of optimization problems SD-2a and SD-2b is to differentiate service classes by routing only. SD-2a fixes the required ratio of mean delays, whereas SD-2b uses classwise cost weights. In Figure 4 the ratio of classwise mean delays of the service classes is presented as a function of the cost weight of the first class.

If we want to differentiate the service classes by routing, the achieved routing obviously differs from the optimal routing, which minimizes the total mean delay $E[D]$. So we use the total mean delay as a function of the ratio of classwise mean delays as a performance metric of service differentiation. In Figure 5 the total mean delay is presented as a function of the ratio of classwise mean delays using

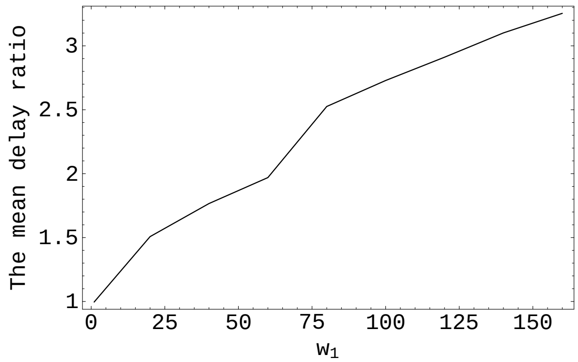

Figure 4: The ratio of classwise mean delays as a function of w_1

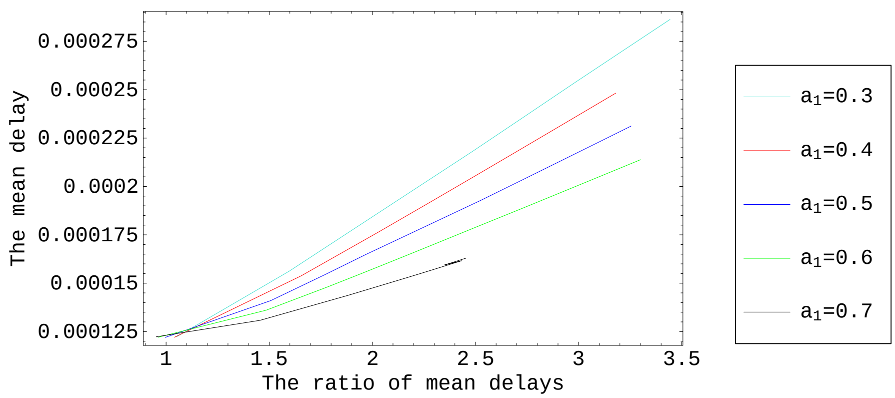

Figure 5: The total mean delay as a function of the ratio of classwise mean delays when a_1 varies between 0.3 and 0.7

different values of a_1. Increasing the relative traffic demand a_1 of the first class reduces the mean delay at any fixed level of service differentiation.

4.4 Service differentiation by scheduling and routing

Finally, we combine two approaches, service differentiation by scheduling and service differentiation by routing. As Figure 6 shows, the mean delay is constant up to the approximative upper bound. After that the mean delay increases almost linearly as greater differentiation is achieved.

In Figure 7 we compare two approaches, service differentiation by routing and service differentiation by scheduling and routing. The left part of the figure shows the classwise mean delays and the right part the total mean delay as a function of the ratio of classwise mean delays. We can see that the mean delay of the class with the higher priority decreases only marginally after the approximative upper bound of ratio of mean delays.

Optimization problem SD-3a differentiates the service classes only at the network level. We like to study the mean delay at the link and stream level as well and compare those to the obtained ratio of mean delays at the network level. In Figure 8 we present a scatter diagram of the mean delay ratios at the link level using different values of a_1. The scatter diagram of the streamwise mean delay ratios is shown in Figure 9. The lines in all figures indicate the achieved ratio of mean delays at the network level. The diagrams show that the variations of linkwise and streamwise delay ratios increase as the ratio of mean delays of the network increases.

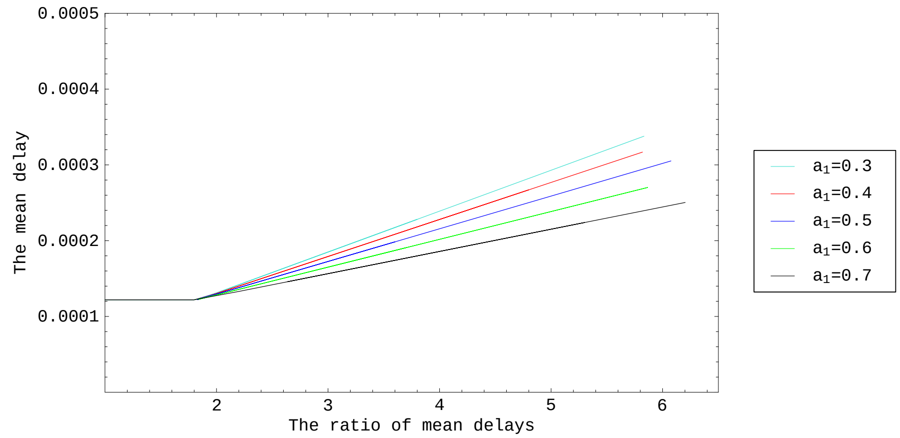

Figure 6: The total mean delay as a function of the ratio of classwise mean delays when a_1 varies between 0.3 and 0.7

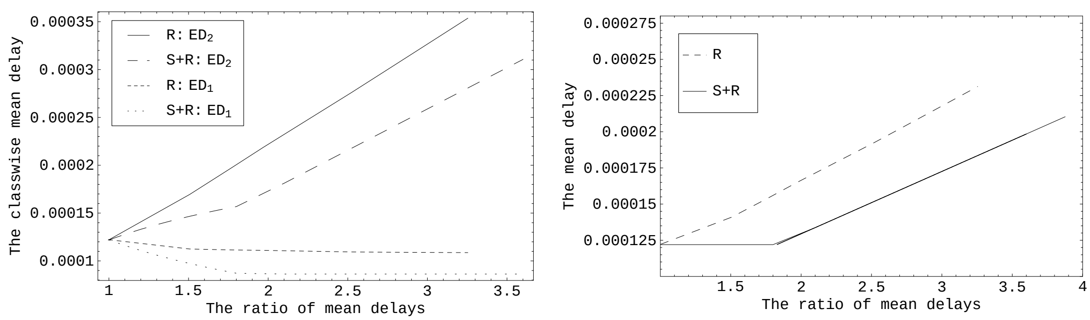

Figure 7: The streamwise mean delays in the left part and the total mean delay in the right part of the figure as a function of the ratio of mean delays when a_1=0.5. "R" refers to the differentiation by routing and "S+R" to the differentiation by scheduling and routing.

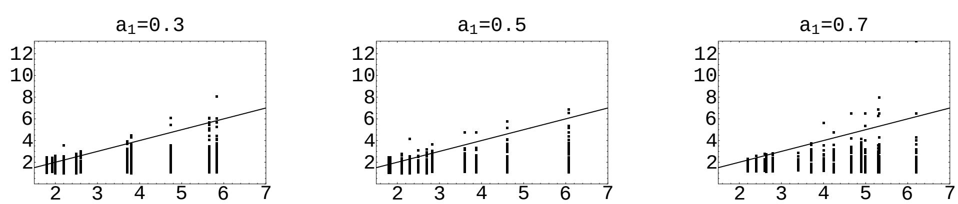

Figure 8: Scatter diagram of the linkwise mean delay ratios as a function of the ratio of mean delays at the network level

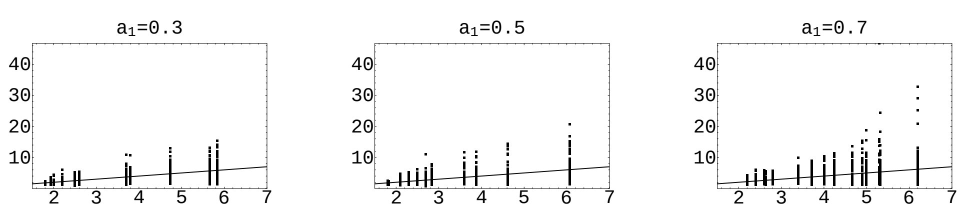

Figure 9: Scatter diagram of the streamwise mean delay ratios as a function of the ratio of mean delay at the network level

5 CONCLUSIONS

MPLS as a new technology brings up new possibilities to enhance the performance of IP networks. The capability to use explicit routes and split traffic along several paths allows load balancing. Using load balancing formulated as an optimization problem as a starting point we have studied how the service classes could be differentiated in terms of the classwise mean packet delay.

We have presented three approaches to differentiate service classes. The first approach relies on WFQ-scheduling. We have presented the approximative upper and lower bounds to the ratio of mean delays achieved just by scheduling. The second approach differentiates service classes by routing only. This means that the service classes with a higher priority should be routed along paths that are not congested. In the third approach we have combined the two approaches.

The optimization problems are implemented and tested in a test-network. In our example, the approximative maximum ratio of mean delays achieved by scheduling methods only is 1.8, which may not be enough. Using routing without and with scheduling provides more flexibility in differentiation but the drawback is the significant increase in the total mean delay.

References

[1] D. Awduche, J. Malcolm, J. Agogbua, M. O'Dell, and J. McManus Requirements for Traffic Engineering over MPLS. IETF RFC 2702 (Sep 1999).

[2] D. Bertsekas and R. Gallager *Data Networks*, 2nd edition. Prentice Hall (1992).

[3] S. Blake, D. Black, M. Carlson, E. Davies, Z. Wang, and W. Weiss An Architecture for Differentiated Services. IETF RFC 2475 (Dec 1998).

[4] R. Braden, D. Clark, and S. Shenker Integrated Services in the Internet Architecture: an Overview. IETF RFC 1633 (Jun 1994).

[5] T. Carpenter, K.R. Krishnan, and D. Shallcross Enhancements to Traffic Engineering for Multi Protocol Label Switching. Proc. ITC-17, Salvador da Bahia, Brazil (Dec 2001) pp. 529-540.

[6] R.G. Gallager A Minimum Delay Routing Algorithm Using Distributed Computation. *IEEE Transactions on Communication* **25**, 1 (Jan 1977) pp. 73-85.

[7] http://www.gams.com.

[8] F.P. Kelly *Reversibility and Stochastic Networks.* John Wiley & Sons (1979).

[9] L. Kleinrock *Queueing Systems. Volume 2: Computer Applications.* John Wiley & Sons (1976).

[10] T. Ott, T. Bogovic, T. Carpenter, K.R. Krishnan, and D. Shallcross Algorithms for Flow Allocation for Multi Protocol Label Switching. Telcordia Technical Memorandum TM-26027 (2001).

[11] A.K. Parekh and R.G. Gallager A Generalized Processor Sharing Approach to Flow Control in Integrated Services Networks: The Single-Node Case. *IEEE/ACM Transactions on Networking* **1**, 3 (Jun 1993) pp. 344-357.

[12] E. Rosen, A. Viswanathan, and R. Callon Multiprotocol Label Switching Architecture. IETF RFC3031 (Jan 2001).

[13] R. Susitaival Load Balancing by MPLS in Differentiated Services Networks. Master's Thesis. Helsinki University of Technology (2002).